

XONA PARTNERS

Hindsight is 20/20: Internet Technologies Trends Retrospect What We Thought And What Happened

Xona Partners

September 2020

As Xona Partners turned 7 in 2020, the year of Covid-19, it's a good time to pause, reflect on the past and plan ahead. A retrospect of some of the most strategic projects we worked on is found in the technology insight papers we wrote, synthesizing our thoughts, views and learning from engaging innovative technology actors and investors with a single sharp focus: advancing deployments of the global Internet Infrastructure.

Many significant trends we saw coming in the mid 2010s became reality. Notable examples include the evolution of network function virtualization, the emergence of the telecom cloud platform, the slow but progressive leverage of Artificial Intelligence by telecom operators as a productivity engine, the use of continuous integration and delivery DevOps framework in telecom service deployments, the deployment of network sharing models and the migration of some telecom services to the edge.

While many of these trends were unclear 5-6 years ago, we bet on them materializing by balancing among the factors that govern the adoption of new technologies. However, we feel it's more important where we bet against certain trends, advising our customers to avoid getting carried away by market hype that falsely projects industry consensus. Some of these include bets against rapid small cell deployments, against a speedy 5G rollouts in mmWave frequency bands, against telecom operators aspiring to dominate the edge cloud and against rapid IoT network deployments. For these trends, we made a point of the inadequacy of business models and lack of maturity in operational models and regulations would prove challenging and would slow down adoption at scale.

The consistency in positioning a vision of the future is born out of our DNA in building our own Internet technology ventures. 20+ years of hands-on development, deployment and market engagement by each of our partners provide us the foundation and confidence to assist our clients define solutions and make difficult decisions.

But we're not stopping here! Our enthusiasm for technology is propelling us forward into new promising fields: low earth orbit space Internet, quantum cyber-security and blockchains applications in Internet business models. Here again, we have been stating our views and defending our claims to help our clients formulate the right approach for their technology and business roadmap. No one can predict the future, in particular the future of technology. But our role remains in bringing a well-founded perspective based on experience, synthesis, and sober analysis of facts to help our clients generate options and navigate the process of decision making.

With this selection of white papers as a retrospect of what we saw coming and what we thought won't happen, we have more confidence to move forward.

The Xona Partners Team
September 2020

Table of Contents

1. **Competing for the Edge: Analysis of Competitive Dynamics Between Cloud Providers and Telcos**
2. **Main Trends in the Edge Cloud Ecosystem**
3. **Robots + ROI: The AI Dimension**
4. **Space Intersects Internet: Opportunities and Challenges**
5. **Emerging Technology Disruptions: Learning from Experiments**
6. **The Critical Dimensions for 5G Fixed Access**
7. **A Foothold in Silicon Valley**
8. **A Perspective on Multi-Access Edge Computing**
9. **Will Open Source Disrupt the Telecom Value Chain?**
10. **RAN Virtualization: Unleashing Opportunities for Market Disruption**
11. **The State and Future of The Home Automation Market**
12. **Internet of Things: Roadmaps and Regulatory Considerations**
13. **Shaping Cellular IoT Connectivity**
14. **Internet of Things – The Turning Wheels of IoT Investments**
15. **Internet of Things – Coming of Age**
16. **Riding The Advance Cloud Deployment Roadmap**
17. **Strategic Advisory – The Case for a Disruptive Model**
18. **Innovations In RF Distribution Networks**
19. **Home based Healthcare & Opportunities for Mobile Operators**
20. **Active Infrastructure Sharing Technology & Business Models**
21. **The Path to 5G Mobile Networks – Gradually Getting There**
22. **Data Sciences Focus – Mobile Eco-systems Contributions**
23. **The Internet Eco-system Value Chain: It's Always Greener on the Other Side. Or is it, Really?**
24. **Online Advertising Real Time Bidding and Opportunities for Mobile Service Providers**
25. **Xona Technology Cycles Deja Vu**

XONNA PARTNERS

Competing for the Edge

Analysis of Competitive Dynamics Between
Cloud Providers and Telcos

Frank Rayal,
Dr. Riad Hartani

July 2020

Edge computing stands at the intersection between cloud providers and telcos, each seeking to carve a role in servicing the enterprise. This raises questions on who will be better able to generate revenue from edge cloud services, and the nature of the competitive landscape between telcos and cloud providers.

To answer these questions, we reviewed the approach of the cloud providers and telcos towards the edge. The cloud providers leverage data centers designed for scalability and efficiency but are physically far from the end user. Migration towards the edge helps them reduce latency and save on the cost of transport to centralized data centers. On the other hand, telcos are in the process of launching 5G networks with the promise of low latency and high bandwidth that can only be realized with edge computing.

The evolution of the edge cloud is a complex topic. Here, we describe an important aspect of this evolution which is governed by many deployment scenarios and applications. Our approach is to segment the market to project the prospects of the cloud providers and telcos. In one dimension, we have the type of cloud: public and/or private cloud. In the other dimension we have on- and off-premise edge computing. We believe these segments cover the most important deployment scenarios required to assess competitive dynamics.

Evolution of the Public Cloud Edge

AWS, Microsoft (Azure) and Google have close to 60% of the public cloud market revenue. They are rapidly developing edge services to cater to their enterprise clients. The first edge solutions focused on device-side applications that benefited from local processing in low bandwidth availability and reliability. Recently, within the past year, AWS and Microsoft released new edge solutions which placed instances of their public cloud infrastructure on enterprise premise (a single or few racks of servers) or at the telco.

The AWS services include Wavelength which hosts infrastructure at the telco central office and Outpost which hosts infrastructure on enterprise premises. Similar services by Azure include Azure Edge Zone with Carrier and Azure Edge Zone for enterprises hosting on premise.

Cloud providers view the edge cloud as an extension of the public cloud. The same tools for automation, deployment and security controls are used in both cases, as are the service application programming interfaces (APIs). Both edge and cloud services run on the same infrastructure and have the same operational consistency for functions such as upgrades, patches and versions. In both cases, applications can scale up or down and are billed based on resource utilization.

AWS announced partnerships with Verizon, Vodafone (UK, Germany), SK Telecom, KDDI. Microsoft announced AT&T, CenturyLink, Etisalat, NTT Communications, Proximus, Rogers, SK Telecom, Telefónica, Telstra, and Vodafone. These are non-exclusive agreements, so operators could sign with different cloud providers, just as cloud providers could sign with different operators. The telcos provide their central offices as hosting locations. The computing infrastructure is tested with the network and optimized to minimize latency.

Evolution of the Private Cloud Edge

Analysts estimate that only 20% of enterprise workloads run on public clouds, leaving the remainder 80% to run on private infrastructure. Private cloud providers are seeking to capitalize on this market by enabling enterprises to implement a hybrid-cloud model where workloads run on the most suitable platform for the desired task, including the edge cloud. This means solutions to meet the different requirements for workload deployments in public cloud, private cloud, virtualized or bare metal; and to allow enterprises to automate provisioning, manage and orchestrate functions across multiple locations.

There are many players in this sector including both established companies and startups typically addressing public-private hybrid clouds. Key players include VMWare, RedHat which is part of IBM, Ubuntu, Volterra and many other players. From a telco perspective, MobileEdgeX is notable for being a spin-out from Deutsche Telekom with a business plan to provide edge cloud PaaS services operating from locations leased from telcos.

Comparative Strengths and Weaknesses of the Cloud Providers

The cloud providers already possess a number of key strengths in the evolution towards the edge which are the following:

1. **Technology and infrastructure:** The infrastructure that forms the cloud – the data centers, software stacks and backbone connectivity – provides a scalable global platform to host enterprise services. The edge is considered an extension to the cloud to allow enterprises run workloads in the most suitable location, and to change that location at any time depending on desired performance and cost. Edge applications can be managed and controlled from the Cloud. The smooth migration from centralized public cloud into the distributed edge is a key advantage of cloud providers. Telcos can provide information about the performance of different locations, such as latency and QoS. But ultimately, the cloud provider owns the platform; and enterprises make their purchasing decisions based on which platform best meets their requirements.
2. **Enterprise client-base:** Cloud providers have an established client base of enterprises for their wide range of services (SaaS, PaaS, etc.). These enterprises could benefit from edge services either to improve the performance of existing applications or to develop new applications. The technology developed by cloud providers took years and billions of dollars to develop. In the meantime, cloud providers perfected their operation and delivery model. In contrast, telcos provide connectivity but few applications and services above that. While there is much room to grow in the enterprise cloud market as many enterprises still rely on private clouds, the telcos are at a competitive disadvantage in winning that business.

3. **Developers:** The cloud providers have a large number of developers building applications on the cloud platforms. Developers can use the same development and management environment for both the edge and cloud services. There is no equivalent ecosystem of developers for telcos, which would be difficult to replicate especially due to the fragmented nature of telcos. Since applications drive revenue, this is one of the most critical aspects. Telcos actually recognize this shortcoming as evident in telco-led ETSI MEC industry group identifying the application ecosystem as a challenge, and the creation of a developer group in the Telecom Infra Project (TIP).
4. **Ecosystem:** Complementing the developer community is the ecosystem that exists around cloud services. Many applications and services are available to accelerate development of new services on public clouds. Telcos would have to replicate that which would again prove challenging given the fragmentation of the telco community.

Despite these strengths, the cloud providers suffer from a major weakness: lack of physical presence at the edge of the network. Cloud providers leverage hyperscaler data centers for scale and cost efficiency. They have also partnered with other data center operators to get closer to users. Yet, they remain far from being integrated into the connectivity network which is necessary to achieve the ultra-low latency and jitter performance.

Comparative Strengths and Weaknesses of Telcos

The key telcos strengths related to the edge are as follow:

1. **Location and physical assets:** Telcos have hundreds, even thousands, of central offices in cities across their service areas. The evolution of central office technology has left many of these locations vacant or with unused space. Some telcos even proceeded to sell some central offices and aggregate operations into a fewer number (e.g. Deutsche Telekom and NTT Docomo). We exclude towers and cell sites because service providers would not be able to capitalize on such assets because: a. Many telcos sold their tower sites to infrastructure companies; b. There is limited space at tower sites for edge computing hardware; and c. The architecture of the mobile network doesn't lend itself to placing edge computing infrastructure at the tower, at least for the time being.
2. **Access to subscribers:** Mobile network operators sell connectivity services to over three billion subscribers. That makes them an ideal channel for cloud providers and over-the-top application providers (OTTs) in B2C model where the end-customer is a subscriber or an IoT device, including drones and future autonomous vehicles.
3. **Access to enterprise:** This is an arguable strength. Telcos with strong fixed access business typically have better access to enterprises than pure mobile

network operators to whom the enterprise is a group of individual subscribers. Some service providers still maintain and operate data centers, especially in markets such as Europe, Japan, the Middle East and other regions.

On the other hand, telcos suffer from a number of weaknesses, such as:

1. Fragmentation and lack of global scale.
2. Lack of understanding in building software and applications at scale.
3. Edge cloud technologies, which are software-based, are not fundamental to telcos' core expertise.
4. Edge cloud services require operational practices that many telcos failed at providing in the past.

In summary, the strengths of the telcos are the weaknesses of the cloud providers and vice versa. This makes a good argument for a synergetic relationship.

The Competitive Landscape

The edge cloud includes different deployment models, such as on- or off- the enterprise premise. On-premise edge implies physically locating the computing, storage and networking infrastructure at the enterprise. Off-premise edge implies locating the edge infrastructure elsewhere, close to the enterprise, but not physically on enterprise premises. To understand the opportunity and dynamics between cloud providers and telcos in the edge cloud, we summarize the analysis in the two tables below first by approaching the edge from a public cloud direction, followed by approaching the edge from a private cloud direction.

Table 1. Edge dynamics from a public cloud approach.

Enterprise On-Premise Edge	Enterprise Off-Premise Edge
<ul style="list-style-type: none"> • An emerging area primarily complementing cloud services and mitigating its shortcomings. • The cloud providers are beginning to offer new edge cloud services as extensions of cloud platforms: e.g. AWS Outpost and Azure Edge Zone. • The cloud providers reduce barriers to adoption by providing the same development, management and operational environment. • Telcos don't have a public cloud play and would be limited to providing connectivity services¹. • Telcos remain limited to providing connectivity services. • Collaboration between telcos and cloud providers benefits both parties. 	<ul style="list-style-type: none"> • Cloud providers dominate public cloud services while telcos don't have such a play. • Cloud providers co-locate instances of their cloud infrastructure in telco central offices transforming them into edge data centers. • Cloud providers leverage the edge as an extension of the cloud while telcos leverage their physical assets and proximity to end-users. • Collaboration between public cloud providers and telcos enables low-latency applications and reduce data transport expenses. • Services such as AWS Wavelength and Azure Edge Zones with Carrier address this market segment. • Close integration with the telco cloud brings further value to cloud providers' services. • The role of telco is primarily providing real estate facilities for the edge data centers. • While telcos could opt to block the cloud providers², telcos would not be able to provide a competing offering.

¹ There are exceptions such as NTT Docomo in Japan.

² As is the case with Alibaba and the service providers in China.

Table 2. Edge dynamics from a private cloud approach.

Enterprise On-Premise Edge	Enterprise Off-Premise Edge
<ul style="list-style-type: none"> • The status quo for the enterprise which owns the edge hardware & software running on its private cloud. • Where the enterprise benefits from 5G for its own use (enterprise network), the enterprise has the choice to own and operate the 5G network, or lease it from a third party that manages the network. • Telcos could provide private wireless networks with a user plane function (UPF) on enterprise premises and play a similar role to a mobile virtual network enabler (MVNE). • Telcos don't yet have such a strategy today (except for trials in Europe). However, such business models would need to consider hybrid cloud models to improve the value proposition for the enterprise. • A cloud provider, such as Microsoft, could provide a hosted core network service. This relegates the role of the telco to a pure connectivity provider. • Hybrid private-public cloud models are evolving to address this market with solutions from the likes of Google and RedHat. This approach could be complementary to telco services. 	<ul style="list-style-type: none"> • A potential opportunity for telcos is to provide hosted edge services in their central office data centers offering enterprises tight integration with the telco network for maximum performance. • Telcos could choose from a few available platforms such as MobileEdgeX, OpenStack or VMWare³. • A telco-hosted public cloud service - e.g. AWS Wavelength and Azure Edge Zones with Carrier - competes with this model, potentially pitting a company like MobileEdgeX or the telco itself against the public cloud providers.

³ System integrators such as WiPro and Infosys are among such players in addition to many cloud providers of the type of RedHat and VMWare.

Summarizing the competitive landscape from, we arrive at the following simplified matrix to describe the interaction between cloud providers and telcos.

Table 3. Competitive landscape between telcos and cloud providers.

	Enterprise On-Premise Edge	Enterprise Off-Premise Edge
Public Cloud	<ul style="list-style-type: none"> A domain for the cloud players where telcos' role is providing connectivity services. 	<ul style="list-style-type: none"> A cooperative partnership between cloud providers who supply the technology and service platforms and telcos who host the cloud providers' infrastructure in telco central offices.
Private Cloud	<ul style="list-style-type: none"> A new market where cloud providers leverage software solutions and telcos leverage connectivity services. Enterprises can opt for hybrid cloud models that play in favor of cloud providers from monetization perspective. 	<ul style="list-style-type: none"> Potential competitive segment between telcos and cloud providers. Telcos could block cloud providers but will need to address telcos' inherent weakness in providing cloud services.

Synergies Between the Cloud Providers and Telcos

Telcos have met successive failures in cloud services, first as public cloud providers then in building their own cloud for their own services. Today, telcos rely on public cloud providers for these services.

While some would position the edge as another opportunity for the telcos to develop a cloud play, our analysis points towards complementary dynamics between telcos and cloud providers. This is particularly the case in relation to consumer services over wireless networks. On the other hand, the enterprise segment could see competitive behavior although both cloud providers and telcos will have to co-exist. We illustrate with two examples:

1. **Complementary coexistence:** AWS Wavelength, Microsoft Azure Edge Zones with Carrier are examples of how telcos and cloud providers could collaborate: Telcos are resellers of cloud services and technology. This helps drive new business to both parties. Telcos leverage their proximity to end-users while cloud providers develop complementary services to their cloud offering.
2. **Competitive coexistence:** Microsoft's acquisition of Affirmed Networks allows it to host a virtual packet core and provide it as a service to connect enterprise radio

nodes in unlicensed (NR-U), shared (CBRS) or enterprise licensed-spectrum. Such a managed service relegates the role of telco to connectivity provider. The telcos would lose a new revenue opportunity for managing the enterprise private wireless network. This has parallels with OTTs services where telcos cannot monetize services beyond connectivity.

Conclusion

The public cloud providers have an advantage over telcos in capitalizing on edge cloud services. This is due to technology, ecosystem and business models. Nevertheless, there are opportunities for telcos because the edge cloud is heterogeneous and the promise of many emerging technologies is yet to materialize. The cloud providers have a head start in technology and operations which creates an uneven playing field tilted to their advantage. The edge cloud is diverse and provides many areas where both cloud providers and telcos could collaborate.

XONNA PARTNERS

Main Trends in the Edge Cloud Ecosystem

Frank Rayal,
Dr. Riad Hartani

January 2020

The decentralization of the Internet through edge computing brings a new set of challenges that require new solutions to meet the performance and cost requirements of the edge cloud. This creates opportunities for investments and M&A across the technology ecosystem as these recent examples indicate:

- Equinix acquires Packet, a developer of bare metal automation platform. Packet received investments from SoftBank Group, Dell Technologies, Capital Battery Ventures, Third Point and Samsung NEXT.
- Siemens acquires Pixeom a developer of Docker container-based solutions to deploy and orchestrate cloud applications on commodity hardware on premises. Siemens plans to use the solution in factory automation.
- Pensando emerged from stealth in October 2019 having raised \$278 million to date. Investors include Cisco, HPE, Lightspeed Ventures, Equinix and Goldman Sachs.
- Volterra which provides a platform for deploying applications in multi-cloud and edge computing environments raises \$50 million in funding from Khosla Venture and Mayfield in addition to other strategic investors.

In this article, we review major trends in sectors fundamental to realizing the edge cloud. For context, we highlight key drivers that stimulate the rise of the edge cloud.

The Edge Cloud Drivers

A few trends are defining the evolution of the edge cloud:

1. Extending successful enterprise cloud services towards the edge. This means the harmonization of technologies, development environments and business models of the cloud with the network edge.
2. Optimizing the cost of data transport between the network edge and the cloud infrastructure. The success of cloud services places increasing demand on transport capacity. The edge cloud optimizes the cost structure of the end-to-end network.
3. Meeting the performance requirements of emerging applications requires placing the compute, storage and networking infrastructure at the network edge. Such applications include for example virtual and augmented reality, robotics and automation, artificial intelligence and machine learning.
4. Regulatory requirements for data localization and consumer privacy rights.

The technology ecosystem approaches the edge cloud on the basis of one or more of the above drivers. For instance, cloud players have an interest in extending cloud services. Telco players have an interest in improving the performance of networks to monetize edge applications. The enterprise has interest in optimizing cost and performance while maintaining compliance with regulatory requirements.

Edge Cloud Ecosystem Evolution and Trends

To give a 360-degree perspective on developments in edge computing we cover a few key sectors that form the foundation of the edge cloud: Cloud players/hyperscalers, data center players, silicon vendors, software stack and hardware (servers and storage).

Cloud Players / Hyperscalers

The public cloud players view the edge cloud as an extension of their services. They seek to reduce reliance on the connectivity layer between the enterprise and the cloud.

Several applications are driving the extension of cloud services to the edge. Important services include data intensive Artificial Intelligence (AI) and Machine Learning (ML) applications in different use cases such as video surveillance and image recognition. They also include IoT applications to scale the deployments of sensors and devices. To meet the requirements of these services, the cloud players provide scaled-down version of their cloud software environment to develop applications that run efficiently at the edge and synchronize with the cloud when possible and desired.

Another trend is placing instances of the cloud infrastructure at local data centers, enterprises of telco central offices to improve performance metrics such as latency and jitter (an example of this is AWS Wavelength service).

Data Center Players

Private data center operators are physically positioned closer to end users and enterprises than the hyperscalers. This allows them to provide edge service to their enterprise customers leveraging their proximity. Additionally, some of these players are leasing part of their facilities to the public cloud operators to locally host instances of their cloud infrastructure (example of this is Equinix and its relationship with Azure).

Another key trend is the evolution of micro data centers that could be as small as around 300 sq. ft. in size. The infrastructure design of these data centers is unique to support high-density computing with power density exceeding 1,500 W / sq. ft. An example of this includes VaporIO and EdgeConneX.

Silicon Vendors

There are two key trends related to silicon for edge computing: 1. Increased variety of types of processing; and 2. Low-power computing and storage.

The 'cloud' is largely powered by general purpose processors based on the x86 architecture. Applications such ML made it necessary to develop different types of engines to handle complex and arithmetic intensive processing. These engines include Graphic Processing Units (GPUs), Field Programmable Gate Arrays (FPGAs), and Tensor Processor Units (TPUs). The variety of processing is giving rise to different types of System on Chip (SoCs) that combine different functions on one chip.

To place computing at the device or enterprise, low-power compute engines and storage become critical to meet field deployment models. Of the many examples in this segment to list, we mention Google TPU and edge TPU solutions targeting machine learning applications.

Cloud and Orchestration Software Stacks

The edge cloud software stack is one of the most critical elements in the overall edge ecosystem. The edge cloud is fundamentally a highly distributed cloud concept that encompasses different types of compute infrastructure including servers in data centers, gateways, and different types of edge devices. This requires software solutions that bridge the centralized cloud with the edge cloud, in addition to different solutions to control and manage the edge cloud.

A key rising area are the cloud-of-cloud solutions that seek to allow enterprises deploy workloads across multiple clouds. Another area includes software to deploy and manage microservices at the edge, including software to manage different types of compute and storage infrastructure.

Enabling the telco edge cloud is an active area for software development as exemplified by many open source projects that address the telco edge cloud, including extension of OpenStack features to meet the edge deployment requirements. Moreover, enabling the telco edge cloud is giving rise to a number of companies that are developing solutions to broker deployment of edge workloads between developers and the fragmented telco virtualization infrastructure.

Additionally, there are a number of open source projects and companies in the process of developing and distributing edge stack for enterprise and IoT applications. An example is ioFog by Edgeworx.

Hardware - Servers & Storage

Hardware for the edge cloud has some unique requirements because of the intended use case and deployment model.

Among the key trends in edge hardware is the integration of different functions into a single unit, for instance compute, storage and networking into a single rack unit. Another trend is the rise of “data-center-in-a-box” solutions where compute, storage, networking and power are packaged into a single enclosure. The computing could consist of different types of processors depending on use case (e.g. x86, ARM, GPU or other). Such solutions have various use cases. For instance, they are used where the cloud could not be reached easily or cost effectively. Initially, such solutions were used for data storage, but increasingly compute processing is being integrated to process data at the edge to the extent required by the application.

Concluding Remarks

The edge cloud is a catalyst for innovation across the entire technology ecosystem. Cloud services have proved to be successful, but requirements for data localization, cost and performance optimization create a valid business case for edge computing services. The edge cloud is necessary to launch and scale many applications such as industrial automation, autonomous vehicles including drones, robotics and IoT. In this paper we reviewed developments in a few important segments. However, many other sectors also play an important role, such as security, networking and distributed ledger technologies. All these will make the edge cloud a key area of investment and M&As in the years to come.

XOONA PARTNERS

Robots + ROI: The AI Dimension

Dr. Anurag Maunder,
Dr. Riad Hartani

September 2019

Context

Robotics have evolved immensely over the years. Also, e-commerce along with scarcity of labor is creating an unprecedented demand for articulated robots, and in particular automated mobile robots. For potential customers, robots bring unprecedented benefits over traditional manufacturing and warehouse operations. However, the challenges associated with the robot deployments are multi-dimensional. Robotic vendors need to not only integrate the mechanical, electrical and other innovations in the latest designs, but also integrate a lot of the Artificial Intelligence developments to take care of exception scenarios that humans handle naturally. AI in this context requires integrating a wide range of machine learning techniques that span from symbolic and logic AI to the various instantiations of numerical AI.

The biggest barrier to Robotic deployment, despite all the advances, is the return on investment (ROI). AI has a potential to address the ROI challenges associated with Robot deployment. It brings in new challenges as well. We discuss the fundamentals of corresponding trade-offs in this paper, building a thesis along with that. We do so, based on our hands on involvement in the design and deployment of robots for specific industrial applications, integrating various AI models within that. Some of our experiences are summarized in this note.

Robots ROI Revisited

ROI, of course, is driven by the robot's purpose. The best way to illustrate the ROI logic is via a real world example, from the ones we have been involved with in terms of design, conception and deployment, and from which we can extrapolate some of the conclusions.

Typically it is common for robotic vendors to claim that they can replace a minimum wage human being. We can use the minimum wage as a baseline for our illustration, from which we explore how introducing AI extensively in the robots, gets factored into the ROI analysis. Within this specific scenario under consideration, the worker-replacement based value proposition alone caps the potential revenue per robot to the max of prevalent minimum wage. To earn that minimum wage the robotic vendor has to ensure:

1. The speed and accuracy of the robot is as high as human beings
2. The capability of the robot includes the ability to handle exceptions for the particular task similar to what a normal human worker would handle.

We believe that the only way the above can be achieved is through an extensive use of AI in Robots. The AI is not just the part of robot's task automation but is an integral piece throughout the operations cycle, from automation to diagnostic to support, among other things.

The first challenge is the speed of the robot which depends on the complexity of the job. As an example, consider a task that requires a human being to move a cup from one place to another. It can easily be generalized to include various types (size, shape, purpose) of cups, glasses, thermos etc. However, this simple generalization may be challenging for a robot without extensive AI capabilities. It will likely fail to pick many types of these

items without sound AI integration. However, introducing AI also potentially reduces the speed of the robot (note: this can be mitigated by high powered GPUs, playing into the cost trade-offs). Depending on the design of AI, the accuracy of robot will be a little sub par as compared to human beings. The reason is the exceptions that occur with day to day mundane tasks. Human beings adjust to the exceptions naturally whereas the robots, while much better than machine in terms of handling exceptions, are nowhere close to human beings. This can primarily be mitigated using extensive AI learning loop.

Following the logic of the same scenario: while it's easy to achieve an accuracy of 95%, reaching 99.99%+ accuracy for robots requires a lot of effort in terms of AI modeling, training and tuning. Adoption of AI in robots is the cornerstone of the ROI argument we would like to bring. The cost/accuracy of AI enabled robots can achieve a commercially viable ROI, justifying deployment at scale.

Building on our robot deployment example, the following scenario is also typical: Given the fact that robots would make mistakes, it is common for customers to deploy an overseer (associate) for robots on the floor. The number of people needed to oversee the robot farm depends on the accuracy and intervention needs of the robots. A good number is to have one person manage 10 robots but often times its 1 person managing 5 robots. This means that the lack of good AI reduces the potential revenue by 20%. Furthermore, mitigating some of the potential robotic halts will require the robotic vendor to remotely manage a robot: reboot, restart or request an onsite person help to get the robot reengaged in the production pipeline. If one person at vendor site manages 10 robots, then the vendor cost for managing the robot fleet increases by 10%. This further highlights the need of comprehensive AI needs for Robotic operation..

From the above example we show that a robot may earn only about 68% of the Minimum wage. Robot Earnings = Minimum wage * 0.85 (low speed and accuracy) * 0.8 (on site help needed) that equates to 68%. Here again, the approach to reduce the human intervention and hence associated costs requires embedding more intelligence in robots. Furthermore, robots performance monitoring would be required to identify and accelerate learning loop, which again, calls for analyzing its respective AI modeling, learning and tuning cost tradeoffs.

There are more overheads in terms of power and bandwidth needs so it's a good thumb rule to assume 50% of the minimum wage per robot. This means a robot working for 16 to 20 hours will typically provide the customer of the robot an equivalent ROI of one person in 8 hours. Currently, we are ignoring all the security needs, such as human access to robot and management of robot introduced contamination / risk, to keep the math simple. The security implications of robotic deployments will be covered in another paper where we discuss how AI can help with security, from denial of service to intrusion detection to software vulnerabilities management.

From the above calculations we can deduce that even if the robot is fully utilized (20 hrs a day, about 250 days a year), it can barely earn \$50K a year with fully loaded minimum

wage of \$20 (Robot will get paid \$10 per hour based on our 50% overhead explained above). Most places, the fully loaded salaries are closer to \$15 so the net revenue is going to be less. This lost revenue due to errors from exception handling, speed loss, operator and associate overhead can be significantly reduced by a properly designed AI solution.

There are in fact more challenges. Usually, no single repetitive task that robot replaces has consistent demand. Generally the tasks are seasonal. For a few months, the warehouse is extremely busy and for other months it's rarely busy. It's fair to assume that Robots will not be utilized about 50% of the time. Thus the take home pay for robot gets further reduced to \$25K in the above example. A general purpose AI based solution can help repurpose the robot for different tasks to mitigate the underutilization scenario.

Given the high upfront cost of the Robot and uneven utilization, it is now common for customers to pay Robot vendor on a pay-per-use model, which is a model that is fast becoming the norm.

If the robot vendor wants to recover the money of hardware in about one year then the cost of the Robot should not exceed \$25K (even with our simple calculation which ignore the costs associated with power, bandwidth, security etc). Depending on the sensors, electronics, arms and mechanicals; the cost of the robot can be much higher. The burden on Robotic vendor (aside from the user of the robot such as manufacturing facility or warehouse) needs to include breakages, operational overhead of monitoring robot by vendors, customer support, training, software upgrades etc. So the cost to Robot vendor is lot higher than just the Bill of Material.

Most of these aspects are today addressed by the inclusion of artificial intelligence models in the lifecycle of the robots design and deployment. All of these AI models call for factoring in associated costs at every level of their design and deployment. Hence the robot vendors are forced to provide high value for very low return. The ROI math makes it tough for robotic vendors to innovate until they can bring the cost down significantly, and having full control on the AI integration dimension becomes primordial.

Robot ROI: Further into the AI dimension

There are however strong mitigation factors that we have not yet discussed, and the macro economics are likely to play in their favor over time, specifically in certain industries and certain regions of the world. These are

1. Increased predictability of Robot
2. Increased Labor Shortage
3. Increasing minimum wages

Robots are always on time. They do their work diligently 20 hours a day. Their performance is also quite uniform throughout the day. That can hardly be said for a human being. It is often common to compare the best of human beings with the average performance of a robot. This is not a correct measure. The real measure of the ROI is to include lot more overhead such as hiring expense, management expense and sustained throughout over

a long time. All of a sudden the robots start looking appealing to the customers, when deployed for specific industrial applications.

Labor shortage in critical areas, or during peak season makes the end-user experience difficult to manage. In this competitive world, manufacturing unit or warehouse operator has no choice but to include robots as part of their operations to provide the experience expected from their consumers.

Furthermore, minimum wages are increasing rapidly while the price of Robots tend to go down as more people adapt Robots.

These aspects will over time, increase the ROI for robots deployments, and in conjunction with progressively embedding more AI functionalities, will strengthen the business case for deployments, in many more new areas.

Hence, net net, while it is true that the robotic industry has yet to overcome the ROI challenge, nevertheless it appears to be closing the gap rapidly. Advancements in AI and its rapid incorporation in robot will drive the robotic revolution for next many years.

Conclusion

We revisited the Robots deployment ROI dimension, via an illustrative use case, which through extrapolation, allows us to firm up the following thesis: Robots need to hit very stringent KPIs to make them economically viable. This has led to various false starts in deploying robots. One of the fundamental answers to achieving these KPIs is the need for embedding more and more advanced AI-enabled designs into robots. The caveat is that it requires clear understanding of AI costs, limitations and trade-offs versus benefits. Over time, economics will likely play into strengthening the ROI for robots, but the AI tradeoffs dimension will remain the cornerstone of such ROI equation. The AI dimension integration into robots is one that we have extensively worked on, and applied in real world scenarios. This will be the focus on another discussion paper.

XOONA PARTNERS

Space Intersects Internet: Opportunities and Challenges

Dr. Riad Hartani

May 2019

Space has always been the last frontier for human kind. The Emergence of the Internet has probably been one of the most disruptive and exciting things of the last few decades. As we enter the 2020s, Space and Internet technologies are converging. Potentially. Global leading technology leaders that consider the Internet evolution, in terms of adoption, affordability, performance and reach, as fundamental to their continuous growth, are pouring 10s of billions into Space Internet technologies at the moment. Exciting, yet risky, times ahead.

Disruptions in the fundamentals of Internet infrastructure architecture and design, and the way it is deployed do not happen often. In fact, things have been mostly progressive and incremental over the last two to three decades, since the major shift from the use of circuit switching technologies to Internet packet switching at scale. This has seen a long but steady evolution from Time Division Multiplexing (TDM) based networks to a family of packet based technologies over time, including Frame Relay, Asynchronous Transfer Mode (ATM) and into Multi-protocol Label Switching (MPLS) and their various instantiations, as well as circuit based voice to IP based voice and other multimedia services. In parallel, various iterations of wireless technologies have been deployed, converging to the 5G cellular networks in early stages of deployment today. This has been complemented by the very rapid growth of the broader ecosystem supporting the Internet evolution, in the forms of large-scale data centers and clouds, software operating systems, and over the top applications. In fact, it is primarily this evolution of Internet connectivity models and underlying technologies that led to the growth of the Internet eco-system, as we know it today.

Few interesting paradigms have been emerging over the last few years, with a potential to impact the internet infrastructure design and deployment of Internet based services, with significant consequences on content delivery models, cloud networks, distributed computing and the economics of over the top applications rollouts. These include aspects such as blockchain and decentralized Internet technologies, quantum communications and low earth orbit (LEO) satellite communication networks. This paper focuses specifically on LEO networks, and mostly addresses the challenges to overcome to ensure their potential success. It provides a glimpse of how the technologies, protocols, standards and mechanisms developed for terrestrial and wireless Internet networks can be leveraged to speed up deployments of LEO based communication networks over the next few years.

Simply put, LEO networks are satellite-based constellations that orbit at altitudes below 1200 miles above the earth surface. These constellations have existed for a while, and numerous ones have been launched in the past, with the Iridium network being the most well known from the late 90s. The novelty is in the fact that these recent networks launches are very much focused on enabling global scale Internet connectivity, bringing in a new era of space based Internet technologies. Pretty much all the major Internet/Cloud providers are working on various aspects of such deployments, including Amazon, Google, and Facebook as well as large scale technology players such as Virgin, SpaceX and Softbank along with some of the existing satellite communication providers already present in the GEO (Geostationary Orbit) and MEO (Medium Earth Orbit), as well as venture capital backed startups, and government funded consortiums in China, Japan, Korea Europe and North America. Most constellations launches are being planned during

the 2020-2025 timeframe, with 10s of billions of dollars being invested. At the same time, this is still a high-risk initiative given the technical and business challenges that need to be solved. As such, this is a high-risk high-return equation, and only time will tell on how it will impact the Internet evolution, global competitiveness and Internet geo-politics matters over the next decade.

The new LEO satellite networks being designed at the moment bring in a whole new set of opportunities, taking advantage of the potential low latency, broad reach and high capacity of such networks. The scale of investments going into such initiatives, primarily from the private sector, adds a significant advantage to their potential. These LEO space networks are being designed with the intent of leveraging the mechanisms designed for terrestrial networks such as those for routing, switching, Quality of Service (QoS), resources management, Software Defined Network control, Virtual Network Functions orchestration, Cyber-security, etc.. Yet, a lot of these mechanisms are far from optimal given the characteristics of LEO space networks, in terms of mobility, terrestrial to space wireless links management, and space-to-space wireless links connectivity. In some cases, these mechanisms need to be highly adapted, and in other cases fully redesigned. In fact, these LEO space networks are in early stages of taking advantage of the internet/wireless networking mechanisms that have been developed, deployed and in some cases abandoned over the last 20+ years.

There is an opportunity to leverage state of the art Internet designs and evolving it optimally to enable the deployment of this new generation of space networks. Below is a non-exhaustive review of some of the key aspects that need to be addressed, both in terms of services offering and technology development fronts. For each one of the dimensions considered, we list some of the aspects that require further work, and could take advantage of the various Internet mechanisms and standards out there, for the specific LEO networks context.

Adapting Internet services and customer application offering over LEO networks

- Adaptation of the Services Level Agreements (SLAs) and Key Performance Indicators (KPIs) definition is required. The IP based services SLAs have primarily been defined with terrestrial networks in mind. Adapting them to LEO satellite networks is a must, as it has a direct impact on traffic management/engineering solutions that need to be put in place on the satellites, coordination between terrestrial and satellites networks, load balancing across space segments, among other things, and this on both data and control planes
- Various services targeted by LEO networks are focused on well known internet services offered by existing terrestrial/wireless networks, such as business centric layer 2 and 3 services, Virtual Private Network Services (VPNs) etc. There are new opportunities for services that would leverage the new cost structure of LEO networks deployments in terms of coverage, bandwidth and latency, as well as the potential

new layer 3 routing topologies that they bring such as global Routing with a reduced number of Autonomous Systems, new peering/transit models, among other things.

- The analysis of new services includes aspects that would piggyback on the deployment of distribute mobile edge computing solutions with highly distributed data centers and clouds, content delivery networks, public safety networks, etc.
- There is an opportunity to revisit the technologies and deployment models of peer to peer (P2P) based networks, and leveraging the characteristics of LEO networks in bringing in new topology models for designing and hosting peers' hierarchies and topologies. It would also be interesting analyze how this would complement the ongoing blockchain lead initiatives for incenting the use of P2P networks at scale and the evolution of file systems distributions.
- The emergence of LEO networks opens up new opportunities for the deployment of global Mobile Virtual Network Operations (MVNO) given the large-scale geographical nature of LEO networks and their underlying economics.
- Multi-media services, including voice and video services delivered directly over LEO networks, call for a rethink of the various mechanisms designed for LTE networks, such as those in the IP Multi-media Systems (IMS), roaming models, and inter-connection architectures.
- The global nature of LEO networks, and the new interconnection models it provides with terrestrial wireline, wireless, submarine and cloud networks, has the potential to significantly change the dynamics of rolling out high speed broadband in rural regions, and in particular in the developing world. It is as such, a clear opportunity for a lot of countries to explore ways of speeding up the implementation of their digital infrastructure strategies.

Adapting Internet Routing and Signaling Protocols Design to LEO networks

- Adaptation of the Internet Gateway Protocols (IGP) and potentially Border Gateways protocols (BGP) for global routing to accommodate wireless links with very specific characteristics (this includes satellite to satellite links, ground to space fixed wireless links, mobile users to space wireless links, etc.), and direct impacts on layer 2 and 3 topology information dissemination, path computation, mapping of demand to paths and load balancing over paths.
- Bringing in the consideration of wireless link characteristics in the measure of QoS metrics and their usage for traffic routing, for the earth to satellite links as well as satellite-to-satellite links.
- As LEO networks get progressively deployed, and given the challenges in addressing their specific predictability, reliability and availability characteristics (weather, capacity limitations, etc.), there is a clear need to build network control models that leverage the potential complementarity of other technologies, including 4G/5G,

Microwave backhaul, submarine networks etc. to ensure end to end SLAs are satisfied with the right economics.

- The Handover models typically deployed in 3GPP 4G/5G networks would need to be adapted for the cases of mobile and high velocity satellites, as they call for different mechanisms to ensure data continuity with the appropriate quality of experience requirements. This is even more the case when dealing with dual network elements mobility scenarios, which includes mobile user terminals and mobile satellites.
- The data-path connectivity protocols, centered around the various layer 2/3 IP/MPLS mechanisms, as well as their corresponding control planes, would benefit from potential adaptations that would make them more optimal when carrying payloads over multi-hop space segments.
- The global reach of LEO networks potentially enables a more rapid adoption of internet based services by a larger number of users in the developing world, of IoT services globally and of peer to peer services. All of them requiring a larger Internet addressing space, and in turn, potentially speeding up the adoption of IPv6 addressing. Benefits could go beyond the expanded addressing space itself, and would include opportunities for evolved routing, QoS and security schemes.

Adapting QoS and Traffic Management Mechanisms to LEO networks

- Data path resources management building on top of existing Transmission Control Protocol (TCP) and User Datagram Protocol (UDP), including their various alternatives developed for existing GEO satellite networks, taking into account aspects of high latency, high loss wireless links, compression, QoS signaling, etc., need to be adapted to LEO networks, as their characteristics are very different than standard GEO space networks.
- The design and dimensioning of oversubscription models over LEO space segments have to be fundamentally adapted compared to the models in use in terrestrial networks, given the specificity of traffic models in terms of network capacity demands the variability of the physical/logical space and ground to space topologies, along with the mechanisms available on the data and control paths for short/mid term traffic/resources management
- For a good number of LEO based services applications, the mechanisms in use in 4G/5G packet core networks, to optimize performance and efficiency, in terms of data-path adaptive and reactive optimization would benefit from adaptations taking into account multi-hop space networks characteristics.

Leveraging NFV, SDN and Operational Systems for the deployment of LEO networks

- LEO Space networks are global and hence there is a need to consider ways of deploying SDN and centralized/distributed network controllers and orchestrators in a way that satisfies latency QoS and security requirements and optimizes the cost of deployment and operations.
- This is also the case for the deployment of Operations Support Systems (OSS) and Business Support Systems (BSS) data models, for data ingest, processing and corresponding actions for the management in the network and orchestration of services.
- As terrestrial networks evolve towards NFV models, there is a clear need to leverage these concepts for the design of LEO satellites, for some of the data path functionalities (e.g. routing, QoS, services adaptations, etc.), while considering the constraints of satellites design (Operating Systems, link/data layers, Power, Upgradability, etc.).
- The interaction between the VNFs and the SDN controllers and orchestrators would have to be revisited to take into account the management requirements of satellites as far as dynamic configurability over global topologies
- Interaction of high OSS/BSS layers with the network layer via orchestrators across domains that have been developed for primarily terrestrial networks need to be adapted to LEO networks, as the various messaging / API models would need to include different set of information models and messaging to map the requirements of the data and control paths

Leveraging state of the art cyber-security mechanisms in LEO networks

- Cyber-security for data and control paths would require a new rethinking to accommodate the characteristics of space segments given the constrained functionalities on the satellites, in terms of ability to process, detect and protect their compute and network resources (versus standard routers in terrestrial networks, with way more powerful capabilities), due to the design constraints being considered on space satellites (power, space, cost, upgradability, support, etc.).
- The aspects that relate to data residency for all aspects of network control and management including aspects such as fault management, performance management, billing, etc. would need to be architected very differently given the global nature of LEO networks, and the increasingly local nature of data residency on a per country/region basis
- Opportunity to leverage new key distribution models, including those of quantum keys distribution (QKD) from satellites in space to enhance end-to-end encryption.

Evolving next generation IoT networks leveraging LEO connectivity

- The recent evolution of IoT connectivity services defined in 3GPP, Low Power Area Networks (LPWA) and others could take advantage of LEO connectivity characteristics as far as enhancing the business cases of deployments, as well as the possibility of offering different type of IoT services in rural/remote areas.
- Complementarity between terrestrial IoT networks and space based connectivity networks, provides a new framework for global service providers to deploy retail/wholesale IoT services at scale
- The IoT gateways and backend architectures in use today would benefit from interfacing with the control and management plane of LEO networks to provide an end to end IoT services deployment and cost/functionalities optimization.

Adapting and evolving technology standards and regulations for LEO networks

- Standard bodies have already been addressing the various regulations required for the deployment of large scale LEO networks. However, various open areas remain under consideration given the global nature of LEO networks, the impact on local regulations on a per-country basis, and the various licensing schemes that need to be adapted
- Standards have also been addressing the aspects that relate to the management of interferences risks with GEO/MEO network as well as the various terrestrial networks. This is likely to be an active area of work as the deployments progress.

Major technology and financial investments are going into the deployment of LEO networks at the moment. There has rarely been so much of a push to experiment, design and launch breakthrough highly complex Internet technologies at scale. It is a race between lead technology players, governments, policy makers that is likely to accentuate over the next few years, given how strategic is the Internet infrastructure for the development of nations and technology corporations competitiveness. Yet, major challenges remain to overcome. This includes both technical and business challenges. The intersection of space and Internet technologies is still in its first phases, with lots of learnings from both sides aiming to enhance the joint value proposition.

In this paper, and building on our own work in the design of space Internet networks, we primarily leveraged our own experiences developing Internet based protocols and deploying Internet based services at scale over the last two decades, with specific views on how they can be leveraged for addressing the challenges of LEO based satellite constellations.

The next few years will likely witness a rapid evolution of these technologies, with a possible significant impact on how Internet services will evolve. A potentially high risk high return equation, where there will likely be few winners and lots of losers. Exciting times ahead in terms of Internet evolution, in a world where Internet is, and will continue to be the cornerstone of the development of nations.

XOONA PARTNERS

Emerging Technology Disruptions Learning from Experiments

Dr. Riad Hartani

4 May 2019

Summary

The Internet infrastructure is evolving into a new phase, bringing in new disruptive technologies such as space networks, quantum computing and blockchain platforms. We provide a brief synthesis of some of our work in these areas with the most important take-aways.

Emerging Technology Disruptions: Learning from Experiments

The recent string of large-scale technology investments over the last few years, mostly led by the cloud/internet players, and in areas as varied as cyber-security, space internet, blockchain, quantum computing and the likes, points to some interesting inflection points in the technology innovation eco-system. It basically highlights the rapid emergence of disruptive technologies, that in essence, build on top of the latest disruptive business and technology cycles that we have witnessed over the last couple of decades, centered around the deployment of Internet and data technologies at scale.

We have been heavily involved in working on these technologies, with some of the lead internet/cloud players. Although most of the work is still at design phases and primarily at experimentation stages, some significant observations are emerging, in terms of what will likely end up being the priority and focus in terms of investments and technology developments, where killer applications are likely to emerge, and what challenges one would need to overcome.

Some of the work we have been doing is described along the five key areas listed below:

- **Internet Intersects Space Technologies:**

The space race is on again, and this time primarily focused on building a new generation of low earth orbits (LEO) satellite constellations, to large-scale Internet broadband delivery. At the heart of it, a simple equation: for the cloud/internet players to keep their business models going, more Internet is needed, and to more people around the world. Breakthrough in space technologies have led to a drastic reduction in the cost of launching satellites, building them and operating them, and as such a new era has opened. We have been actively working on some of the latest designs bringing in Internet knowhow into the new generation space and satellite technologies. In fact, a lot of the last two decades of learning deploying the backbone of the Internet infrastructure (4G/5G, Hyper Scale data centers, Submarine networks, Internet wide routing and quality of services, network wide cyber-security, etc.) provide a first set of solutions to a lot of the challenges of Low Earth Orbit satellite constellations, augmenting satellite networks with designs that have been proven and deployed at scale in the Internet.

- **Into the Quantum Era:**

Compute technologies are the common denominator for the growth of a lot of the technologies we are witnessing today, from AI to IoT to Blockchain and others. Evolving them is the challenge to crack for those that would want to win the technology race, and quantum computing has been one of the key breakthroughs to go after, and a lot

of progress has been made in building the new generation of Quantum Computers. Yet, the first major commercial breakthrough of Quantum technologies are emerging in an adjacent area – that of Quantum Internet cyber-security, with the goal of making the encryption technologies way more robust, leveraging a new generation of key distribution and management technologies. Some of these quantum technologies are likely to become the main focus of high security networks, and possibly mandated over time. Our work in this area has been focused on ways of operationalizing these technologies and taking them into the real world, and learnings from the field are so far, very exciting.

- **Blockchain moving ahead:**

Blockchain is by now, a technology that everyone knows about and very few have managed to commercially leverage at scale, and this is not because of lack of trying. Tons of applications are running, and some even commercial, especially in areas that have to do with bringing a new generation of Fintech applications to market. Yet, a lot remains to be done, on some of the most fundamental aspects of it, as far as making the blockchain platforms robust, scalable, usable and manageable at scale. Efforts are full speed into that, but will take few additional cycles on the engineering development side, and results are likely to be seen in the emergence of key breakthrough in the decentralized data management, data sharing and exponentially more efficient use of compute/storage/networking resources at scale. This is likely to be supported by the major initiatives launched by cloud players offering blockchain platforms as a service, leveraging the scale, cost dynamics and ecosystem pull of large scale public clouds. We have been working on the intersection of blockchain and the Internet infrastructure, which from experience will open up a new wave of applications, leveraging these platforms. From there, killer apps will very likely emerge. We just don't them yet.

- **Artificial intelligence itself needing disruptions:**

AI has had lots of lives, and we are just witnessing one of its best times. A new era of computing, the flood of data coming out of all the new Internet business models, and the highly competitive data driven economy, lead to incredible advances in how AI is used and is by now, almost a feature in a lot of advanced products coming to market. Yet, this has been the case for numerical AI specifically, in the form of machine and deep learning models, while the other branch of AI, symbolic AI, has seen very little progress. Our work has focused on developing models where symbolic AI would come in to address some of the challenges of numerical AI, as far as cost of training, complexity of learning and efficiency of reasoning. This in some sense is a repeat of some of the initiatives run in the mid 90s when numerical and symbolic AI converged, and as such, we shall expect a revival of hybrid models over the next few years. This leads us to believe that the next decade will see a lot more of synergies between the different intelligent computing technologies, with AI being one of the most fundamental components, with a new set of applications emerging out of that.

- A new era in the delivery of Web Scale Software:

Approaches for building software systems have changed drastically over the last few years, and at the heart of it, two fundamental drivers: the move to the cloud and the emergence of large-scale open source software and developer communities. A lot is tried, and some is adopted, and becomes the norm. We have seen that with the first generation of cloud based software, leveraging virtualization and cloud compute models, followed by a new era of containerization of software at scale, and into new models showing promise in areas such as server less compute and other models. Yet, a lot of these developments all call for a common thread: the automation of software delivery, and deployment at scale, leveraging advanced API models, machine learning for software integration and delivery, and allowing through that the development of rapid release of software applications at scale. This is bound to continue, and will be a key competitive differentiator for the application developers aiming at leveraging the new generation of cloud based compute architectures.

Besides the ongoing disruptions we are seeing at the moment, one shall expect the emergence of a totally new set of applications and business models over the next decade, that would drastically change what we know today. This is likely to lead to the emergence of new technology leaders over the next decade, displacing the ones we know and live with today. This time, as it was the always the case before, those embracing change would be the leaders to stay, and others would be absorbed or disappear. Put simply, just like basic genetics!

XONNA PARTNERS

The Critical Dimensions for 5G Fixed Access

A techno-economic analysis of millimeter
wave networks

October 18, 2017

Executive Summary

To assess the commercial potential of millimeter wave fixed access technologies, we developed techno-economic models to validate various business case scenarios. This is only one of multiple factors that impact commercial success, but it is important as it focuses the spotlight on critical aspects such as service plans and pricing, deployment process, equipment features and capabilities, spectrum, and ecosystem development.

Fixed wireless access in millimeter wave frequencies emerged as a principal application of 5G technology driven by the business plan of a few service providers. The process of standardizing the technology is well underway and several trials have been completed or are currently underway by leading vendors and service providers. The two leading US operators, AT&T and Verizon, competed strongly to acquire spectrum in the 28 and 39 GHz bands. Verizon has further engaged in 11 market trials to characterize the technology and assess its feasibility. All this has heightened the interest of financial investors and wireless ecosystem players in the commercial potential of millimeter wave fixed wireless access.

Validity of the business case is critically dependent on the number of connected houses per site. There exists a threshold below which the business case becomes highly sensitive to other parameters that quickly makes it unviable, especially in the presence of other competing technologies. In our case analysis, this threshold is 32 houses per cell site. The other parameters include the cost of site lease, backhaul, and customer premise equipment and installation.

The number of connected houses per cell site is directly correlated to the coverage capabilities of millimeter wave technology. Coverage is tightly coupled with the deployment scenario and the capabilities of the equipment. It is crucial to understand the true performance possibilities of this technology, and how it applies to different markets. This understanding helps to guide the feature design required to realize the successful business models.

The success of millimeter wave is largely predicated on the ability of the service provider to acquire the right site location where capital and operational costs could be amortized over a large enough client base. This and other related factors lead us to conclude that millimeter wave access is a niche application that will take longer than current industry expectation to fully materialize as a significant commercial opportunity.

Table of Contents

Executive Summary	3
Introduction	3
Methodology Summary	4
A Market Perspective	4
The Performance of Millimeter Waves	5
The Critical Elements of the Business Case	6
Connected Houses per Cell	11
Cost of Transport	12
Pole Lease Expenses	12
Cost of CPE	13
CPE Installation	13
Key Takeaways	13
Acronyms	13
About Xona Partners	14

Introduction

Fixed wireless access has a challenging business case. There have been many unsuccessful ventures: LMDS/LMCS in the mid-1990s and WiMAX in the 2000's are prominent examples. Unlike previous attempts, the drive for fixed wireless access is now happening from within the mobile ecosystem, driven by large service providers focusing on millimeter wave spectrum (mmWave). This raises questions on market viability by ecosystem players looking to develop products and solutions: How big is the market? Is the millimeter wave market a niche market? And, should we invest in the fixed access market?

Having experienced previous industry cycles, we at Xona Partners learned to pay close attention to the critical aspects that will allow a technology to gain traction and lead to a thriving market. To address questions related to mmWave networks, we developed techno-economic models that tightly represent the business and usage cases. Technology and business aspects are both critical in such an analysis: the technology performance and market conditions must be appropriately modelled to ensure accuracy. In this paper, we outline key factors impacting the business case, focusing on mmWave technologies under the 5G banner. Our target audience is the investor community, both financial investors and technologists looking to invest in mmWave solutions or networks.

Methodology Summary

We leveraged techno-financial models that Xona Partners have developed and optimized over multiple use cases to determine the most critical parameters that affect the business case. The models combine technical performance in select deployment scenarios with financial metrics that allow us to gauge sensitivity of the business case to different technical, commercial and market parameters.

The simulation engine allows us to cost-out the deployment for different applications. The cost model includes the end-to-end network: access, core and transport networks (Table 1). In this paper, we focus on a deployment scenario in typical suburban area in a US city (Figure 1). The deployment scenario features base stations of small form factor mounted on short poles of 10 – 20 meters in height, in residential areas (i.e. mmWave small cells).

The success benchmark in the business case we present in this paper is the number of months to breakeven. To simplify the presentation and focus on key drivers, we had to consider a subset of all operational aspect of a fixed wireless venture. We therefore left out some parameters, such as customer acquisition costs, while understanding their impact on the business case. In effect, the actual breakeven point for a commercial venture would be longer than the value we present in this paper as our analysis presents a ceiling below which actual operating parameters must remain.

Table 1: Capital and operational mmWave network expenses.

	Capital Expenditures	Operational Expenditures
Radio access	<ul style="list-style-type: none"> • mmWave access nodes • Site acquisition, permitting • Installation, test and commissioning • Radio planning & design • Project management • CPEs • Spares • Spectrum 	<ul style="list-style-type: none"> • Site lease • Transport • Power • CPE installation services • Operation and maintenance • Warranties and vendor support
Core Network	<ul style="list-style-type: none"> • Core network elements (AAA, OAM, billing, DHCP, Firewall, OSS/BSS, etc.) • Design services 	<ul style="list-style-type: none"> • Vendor licensing expenses • Operation and maintenance



Figure 1: Example of North American suburban area.

A Market Perspective

We focus our analysis in the 28 GHz band. By strict definition, mmWave implies frequencies between 30 – 300 GHz, however in the present industry context frequencies in the 24 and 28 GHz are also referred to as mmWave. A few operators are heading the demand, analysis and market trials of mmWave solutions, including Verizon, AT&T, NTT Docomo, SKT, and KT. The US, Korea, and Japan are the current market leaders in setting requirements and in planning for mmWave networks – they have a combined population of 500 million. Other markets, most notably Europe, China and India have been relatively absent, with a few exceptions. The US operators are focusing solely on the fixed use case, whereas the Asian operators have been additionally investigating the mobile use case. The leading fixed access application is fiber extension to provide cable, TV and data services. The geographic concentration of interest in mmWave is important for benchmarking potential economies of scale, especially that related to the cost of the subscriber device (CPE), where volumes are necessary to achieve low price.

The Performance of Millimeter Waves

mmWave has significant throughput performance with a channel size of up to 900 MHz. The challenge resides in the coverage performance. mmWaves have limited non-line-of-sight range due to high penetration loss through walls and foliage, and poor diffraction capabilities around obstacles such as rooftops (Table 2). mmWaves are also susceptible to environmental elements such as rain, snow, and sand, which are accounted for during the planning stage. Bouncing signals, signals that come from any direction, are a practical challenge: the strongest signal is not necessarily the one directly from the transmitter.

Table 2: Range performance for system operating in 28 GHz.

	Coverage range at 100 Mbps cell edge throughput		Coverage range at 1 Gbps cell edge throughput	
	LOS	NLOS	LOS	NOS
Outdoor-to-outdoor		354		219
Outdoor-to-Indoor (standard multi-pane glass)	1260	128	428	66
Outdoor-to-Indoor (IRR glass)	140	37	51	18

The combination of the above challenges leads to high performance variability, which translate into the following practical aspects:

- a. mmWave modems cannot be placed anywhere. Rather, they must be window-mounted and facing the base station. Reflective window coating is a hindrance that leads to outdoor CPEs being required.
- b. Few houses would be covered by a cell site in non-line of sight requiring outdoor CPE deployments to improve the range of coverage and offered throughput.
- c. Outdoor CPE installations require truck rolls by installation specialists.
- d. Beamforming technology is necessary to compensate for performance shortcomings, which adds cost of the base station equipment or the CPE, or both.

The Critical Elements of the Business Case

Of the many elements that impact the success of fixed wireless access deployments, profitability is generally linked to only a few key parameters. To explore the impact of some of these parameters on the business case, we take a scenario of a suburban market served by two competing service providers. mmWave systems are mounted on pole in a 4-sector configuration to serve houses 360-degrees around the pole. The chance of achieving line-of-sight connection to a mmWave modem is 50%.

Connected Houses Per Cell

The number of connected houses, or subscribers, supported by a cell site affects how quickly the service provider can break even on their infrastructure costs. This, along with the service price, affects the revenue side of the business case. But unlike the service price, which is bound by the type of service offered and competitive alternatives, the coverage performance of mmWave technology determines the number of served and connected houses. For instance, a larger cell size covers more houses and spreads costs over a larger client set. This is critical as it sets the foundation of the business case and acts as a bias or anchor around which other parameters can be optimized.

Fixed access networks are typically rolled out selectively, targeting certain areas of interest to the service provider. This is advantageous in controlling cost but also restricts one from leveraging economies of scale.

It is advantageous to the service provider to deploy high poles to extend radio coverage. However, residential areas are very sensitive to cell siting. Often, it is not possible to obtain cell site locations, and when a site is secured, the height of the pole is restricted to below 15 m. This is just above the tree line in many neighborhoods, and restricts reach to the first tier of houses around the cell site.

In our deployment scenario, the business case become valid at near 8 subscribers per sector, or 32 per cell site, based on a 4-sectored configuration. A number below that makes the business case unprofitable, as it has high sensitivity to variations in other parameters. As the number of subscribers increases, the business case becomes more robust to other cost parameters.

To showcase the sensitivity of the business case to other parameters, we set the number of subscribers to 8 per sector.




Figure 2: Impact of number of subscribers per cell on the business case.

Based on the coverage characteristics of mmWave and for practical and commercial reasons, service providers will roll out of mmWave networks in areas where they can serve a large number of subscribers, unhindered by municipal cell siting restrictions and physical coverage obstacles. The fixed wireless access use case is selective.

Cost of Transport

mmWave fixed access networks typically require fiber backhaul, as wireless become limited for multiple reasons. In our scenario, we consider mmWave technology being used for fiber extension, hence, fiber is readily available for backhaul. This has a major impact on the viability of the business model. In fact, leasing fiber backhaul for fixed wireless access is highly unlikely to yield a positive business case. Additionally, the service provider will need to control its own transport network expenses.

In our scenario, the cost of transport to the service provider cannot exceed \$550/sector/month (Figure 3A). After that point, the business case would break even in 56 months. Ideally, the cost of transport should be below \$300/sector/month to present a positive value proposition.

Pole Lease Expenses

mmWave fixed access base stations are deployed on poles similar to small cells. The permitting process has proved to be expensive and challenging. Site leases have also shown to be a major roadblock in this deployment scenario, where in many instances costs of \$1,000/month or more are not uncommon. In the case of fixed wireless access, the business case is sensitive to this parameter considering that relatively few subscribers are served by a site, amortizing the lease expense and justifying the value proposition.

In our scenario, monthly expense for pole lease must be below \$150/month (Figure 3B).

Cost of CPE

The cost of the CPE presents a challenge, because it is typically overlooked in the business case, while on the other side, the industry knows that the success of fixed wireless access is predicated on low cost CPEs. This was understood well after high cost CPEs was a driving reason behind the failed LMDS/LMCS technology. Since that time, fixed wireless access proponents either attempted at creating volume through standardization and ecosystem development (e.g. WiMAX), or adapting other massively deployed technologies for fixed wireless access, such as Wi-Fi and CDMA (WLL).

Complexity and low volumes are detrimental to achieving at a low-cost CPE. mmWave technologies typically maintain higher complexity through technologies such as beamforming, to save on other expenses such as truck rolls. It becomes critical for large markets to adopting mmWave technology in high volumes to achieve the cost objectives. In the world of telecom, the volumes range in the millions of SoCs.

In our scenario, the cost of the CPE should remain below \$350/unit. The business case begins to deteriorate quickly above \$550/unit (Figure 3C). With these figures in mind, we could work our way to estimating a detailed BOM cost for a CPE, including the silicon and antenna subsystems.

CPE Installation

Truck rolls are expensive: they require trained teams, equipment, and coordination to fulfill on their mandate. It is the objective of any fixed wireless access technology to eliminate or reduce to a minimum truck rolls to install CPEs. This often led to sophisticated technology incorporated at both the base station and CPE. While additional expenses at the base station could be tolerated, that at the CPE is more pressing. From this perspective, the cost of truck rolls is a complementary cost to that of the CPE.

In our scenario, where 50% of the CPEs will require truck roll, the cost per truck roll should not exceed \$400/CPE (Figure 3D). In the event that more CPEs will require truck rolls, the cost per truck roll must decrease accordingly.

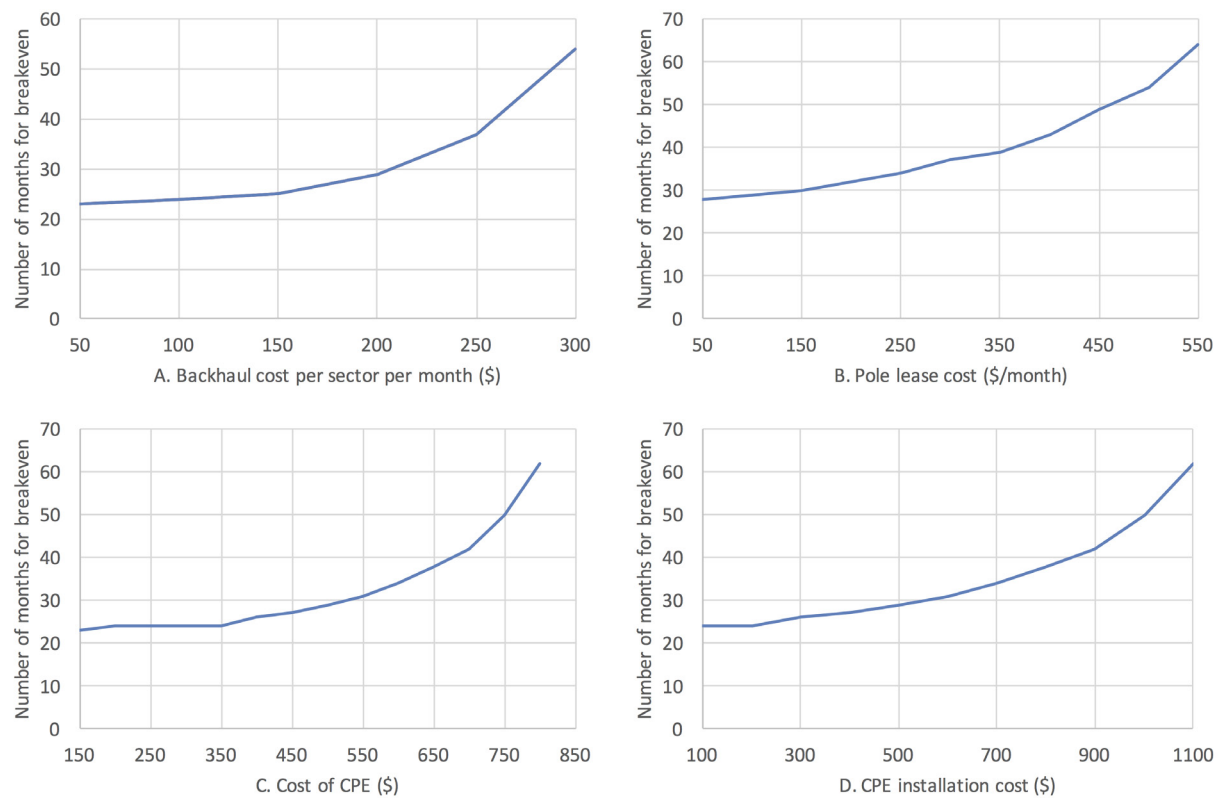


Figure 3: Sensitivity of mmWave business case to key parameters.

Key Takeaways

The selective nature of mmWave renders the technology to niche applications. Requirements for backhaul and cell siting allows a limited number of service providers to take advantage of the technology. These include both wireless and fixed access service providers with fiber assets.

Coupling these conclusions with spectrum availability – an issue that we did not address in this paper, but is of vital importance to achieve economies of scale – lead us to conclude that mmWave will remain a niche technology and will take longer than currently expected to mature and develop: we expect a limited ecosystem for access solutions and a long deployment ramp.

Acronyms

AAA	Authentication, authorization, and accounting
BOM	Bill of material
BSS	Business support systems
CPE	Customer premises equipment
DHCP	Dynamic Host Configuration Protocol
IRR	Infrared reflective
LMCS	Local multipoint communication system
LMDS	Local multipoint distribution system
LOS	Line of sight
mmWave	Millimeter wave
NLOS	Non-line of sight
OAM	Operations, Administration, and Maintenance
OSS	Operations support systems
SoC	System on chip
WLL	Wireless local loop

XONNA PARTNERS

A Foothold in Silicon Valley One (Good) Way to Get There

Dr. Dean Sirovica, Dr. Riad Hartani,
Dr. James Shanahan, Dan Cauchy

June 2017

Summary

This whitepaper should be of interest to any company or government entity that wishes to tap into the Silicon Valley model of innovation and technology disruption. We describe the different models, benefits, and pitfalls of planting a foothold in Silicon Valley. We conclude by describing various execution models to help in this endeavor and maximize the chances of success while avoiding the pitfalls others have experienced.

Introduction

Silicon Valley has become a world-unique and proven birthing ground for disruptive technology startups. This is due to the complex ecosystem at the confluence of University Research, Innovation Spirit, and Venture Capital. This ecosystem is further supported by a large number of businesses and institutions that feed into this ecosystem.

Various players around the world, being corporations, governments, or investment houses have been looking at ways to benefit from the Silicon Valley ecosystem by plugging into it. This is likely to remain the case, and probably, even accelerate. The benefits range from learning and adopting the innovation force of this unique ecosystem, to leveraging it by acquiring new technologies of strategic interest, or to seek exposure to Silicon Valley Venture Capital investment returns.

In this paper, we highlight our learnings and experiences from operating in this ecosystem for several decades, and how this might be applied to benefit other companies desiring a level of exposure to the Silicon Valley ecosystem. The aim is to facilitate a low risk, strategically aligned, presence in Silicon Valley and build an adequate evolution strategy from there.

Our team at Xona Partners can be the gateway platform that would provide a cost-effective foothold in Silicon Valley that best matches the strategic objectives of the parties desiring to benefit from it.

The Drive to Plug Into the Silicon Valley Eco-System

The story of Silicon Valley has been well documented. It started with the defense industry in the 50s and 60s, followed by Integrated Circuits, Personal Computers, the Internet, etc. However another, less visible but significant transformation occurred. Since the 1980s, the US industry has witnessed a shift from in-house innovation (eg: Bell Labs) to inorganic technology acquisition (Venture Capital ecosystem and M&A) as the better model for technology and new business development.

The early decades of Silicon Valley were characterized by waves of innovation in specific industries. Today in Silicon Valley, we see overlapping innovation waves in many industries. These waves of innovation often create synergies that further accelerate innovation and disruption. A good example is the collision of Internet and automotive innovation behind Tesla and the Google self-driving cars. This trend is likely to continue, and put even further pressure on the various global stakeholders in the innovation eco-system to tightly work, integrate and synergize with what's happening in Silicon Valley.

It is our belief that this model of technology and business innovation in Silicon Valley is here to stay despite periodic turmoil in financial markets and the broader economy. Most leading companies will sooner or later have a desire to establish a presence in Silicon Valley in order to tap into this source of innovation and disruption.

Many other parts of the US and the world are trying to emulate Silicon Valley. The numbers speak for themselves. Silicon Valley remains by far the leader in the number of startups and the capital invested in them. In our view, replicating this model in a different geography is not the optimal approach (as evidenced by the many attempts over the years, and we still have one Silicon Valley), and we would argue more for a learn and adapt to context, based on specifics of the local environment, which is a model being successfully pursued by various technology hubs around the world. In that context, we propose that a “bridge to Silicon Valley” is still needed to ensure cross-fertilization of ideas, de-duplication of effort, and adequate access to venture capital.

External Innovation Model

In the US, Venture capitalists invested \$58.6 billion in 4,520 deals in 2016, according to the MoneyTree Report by PricewaterhouseCoopers LLP and the National Venture Capital Association. Although 2016 saw a decline in deals it still represents growth of VC funding when averaged over a few years. Silicon Valley accounts for about half of all US VC deals.

This level of capital fueling innovation ensures a strong supply of talent eager to develop their own ventures. Large corporations find it hard to retain and motivate top young talent. Silicon Valley is full of serial entrepreneurs. In fact, these are typically free-spirited individuals who excel in startups and do not wish to settle in a corporate environment.

This and other factors have changed the old in-house innovation model to one where most disruptive innovation is created outside of large corporations. Corporations are forced to acquire new technologies and new businesses through M&A and partnerships.

With so much activity in Silicon Valley most leading companies are opening offices there to tap into the flow of innovation. This gives them much valued insight and early warning of changes on the radar. There are many examples where companies failed to see the emergence of a significant competitor, especially in industries that Silicon Valley was not known for: BMW was blindsided by Tesla, and Honeywell by Nest. More towards the traditional core of Silicon Valley is the transformation of Cloud Computing, Genetic Engineering, and the Internet of Things. All of these new waves have incumbents scrambling to ride those waves of change rather than be swept by them.

Many corporates have opted to join the Silicon Valley innovation model by opening Startup Incubators, R&D outposts, Corporate venture Funds, and scouting for technologies for acquisition. Some corporates have opted to Spin-Out internal R&D projects into Silicon Valley so they can develop unhindered by the mothership but with an option to be re-acquired at a later time.

Changing Innovation Vehicles

The startup innovation model is rapidly changing. The biggest changes are at the early stages of the startup lifecycle. The startup exit, IPO and M&A, are mostly unchanged (with the exception of the IPO changes caused by the Sarbanes–Oxley Act of 2002). The mid stage VC funding is also relatively stable and well understood. However, the early stages have seen significant changes due to an increasing focus on this early stage by institutional and corporate investors, online versions of syndicate of angel investors, incubators and accelerators, as well as regulatory changes such as Crowdfunding.

Not so many years ago, the life of a startup before VC funding was a very opaque endeavor. There were very little formal statistics gathering or institutional attention. As it came to the forefront, that startups are driving innovation and major business disruptions, investors and corporations have increasingly focused on the earlier stages of startups formation. More recently, the US regulator has made changes that allow for new funding models such as crowd funding for startups. These changes have created a number of early-stage vehicles to stimulate startup creation and early-stage growth.

On the investor side, we have several types of Angel investors from individuals to professionally run angel groups, as well as many early stage boutique VC firms. Most universities have established spin-out centers to facilitate commercialization of the IP generated by their research. There are an increasing number of Incubators, from for-profit, via corporate incubators, to sponsored, and local government supported incubators. There are many startup competitions where winners often get funding and other support. There is an increasing on-line activity that blends social networking with investing to create crowd funding for startups. In addition, there are a number of loosely defined often sponsor supported spaces/venues where entrepreneurs meet, socialize, and work to create startups.

It can be a daunting exercise to understand, track, and engage with this dynamic ecosystem. Inevitably there are many cases of ineffective engagements with Silicon Valley, failed investments, and cost overruns. However, there are also many examples of highly successful engagements, rewarding investments, and lifesaving business transformations.

Planting a foothold in Silicon Valley

A detailed look at the Silicon Valley offices of the many companies present there reveals significant variations in function, mandate, scope, size, and structure. Furthermore, these factors often change within each company over time. This is strong evidence that the winning formula for an effective engagement with Silicon Valley is elusive.

The functions performed by these outposts, is based on various models, depending on goals and strategies. It would include some or all of the aspects below:

- Technology scouting
- Partnerships
- Startup investments from seed stage to mezzanine financing

- R&D
- Incubation
- Due diligence
- VC fund management
- Spin-out & spin-in
- M&A support
- PR & Branding
- Executive education through immersion in Silicon Valley activities

How We Are Approaching the “Silicon Valley Foothold”

Xona Partners has a long experience in Silicon Valley. Our partners have held roles in successful startups, Venture Capital and M&A firms, major tech companies, and led the Silicon Valley offices of global multinationals. Xona has rich relationships and deep networks in Silicon Valley that span decades.

Depending on the strategic needs of our partner, we typically craft a white-labeled presence in Silicon Valley. If the partner wishes to have its own longer term presence in Silicon Valley, we can design an “instant start” Silicon Valley Office for the partner by transitioning from white-labeled Xona staff to partner’s staff through hiring and training in a smooth transparent process without business disruption.

Our aim is to provide the most expedient and efficient way for our partner to establish a foothold in Silicon Valley and reap the benefits that can provide.

In our experience, we have seen the tremendous power of participating at the leading edge of technological and business disruptions. We can confidently predict that, when executed correctly, the investment into a presence in Silicon Valley will have a much higher ROI than the company’s own business. Furthermore, the reduction of the probability of being blindsided, and/or the incubation of a new area of business could have lifesaving consequences. Finally, the executive who has the wisdom to plant a successful foothold in Silicon Valley for their company often receives long-lasting praise

Executing on the “Silicon Valley Foothold” - Partnership Model

A two-phase analysis approach is typically considered:

Phase 1: Scope, Develop, and Deliver a Ready-to-Execute Proposal for the partner’s presence in Silicon Valley

Tasks:

1. Understand the partner’s business objectives, long term strategic drivers, and any existing ideas on how a presence in Silicon Valley can benefit the partner
2. Develop the Strategic Benefits statement and get the partner buy-in

3. Develop the Modus Operandi for the partner's Silicon Valley activities
4. Identify a senior champion for this project within the partner's organization
5. Develop and deliver to the partner the Ready-to-Execute Proposal for their initial foothold in Silicon Valley and evolution thereof

Phase2: Develop the partner's Silicon Valley presence and local engagement model

This phase will be the implementation of the Modus Operandi defined in the design phase. Some examples of the Modus Operandi might be as follows. We shall note that this is very customizable to the partner and may include some mixture of all the examples below including any additional activities defined prior:

1. Scouting

Scout for startups and activities of strategic interest to the partner. Develop "landscape analysis" and deliver to the partner. Facilitate direct-engagement of ecosystem players with the partner.

2. Stimulate Innovation

Depending on the partner needs, organize workshops, events, hackathons, etc. Engage with and stimulate University research groups. Develop "innovation training" seminars for partner's business units.

3. Take equity for option value

Making an equity investment in a startup can give a partner valuable options down the road. These include (a) unique intelligence and visibility into that ecosystem, (b) an option to steer the direction of the startup and technology development, (c) an option to acquire the startup and/or prevent it from being acquired by a competitor.

4. Take equity for investment returns

The returns from Venture Capital can be very attractive. Many investment firms choose to allocate a portion of their portfolio to Venture Capital. We can provide access to a large and diverse early-stage deal flow and tailor the investments to the partner's objectives.

5. Spin-out

Many R&D projects get stifled in a corporate environment spinning them out into the startup ecosystem could significantly improve their chances of success and adding value to the partner. We can facilitate the spin-out process to ensure successful launch of the entity. Spin-outs can also be used to divest the partner from product lines or businesses that are no longer strategic for the partner. This is a way to extract value from an activity that would otherwise just die or distract from the core strategic direction. We can assist in finding a buyer and/or launching the activity as a standalone company with potentially adding external investments if needed.

Select Illustrative Completed Case Studies:

Three distinct use cases, out of the various the Xona Partners team has conducted, have been selected to illustrate the “foothold in Silicon Valley” execution model

Case Study 1: a proxy innovation and venture capital arm in Silicon Valley

A worldwide Internet and Telecom Technology leader decided to place a foothold in Silicon Valley. The objective was three-fold:

1. To tap into and absorb emerging innovative technologies that have significant disruptive potential in telecoms and IT.
2. To strengthen the engagement with startups with a venture capital component that creates additional leverage and de-risks the future.
3. To act as an early warning radar to changes in macro business models caused by technology disruptions.

The Silicon Valley office was structured to best accomplish those objectives. It had very strong executive and working level relationships with HQ in Europe. It was mission critical to maintain a good connection with the core business units and engage them in Silicon Valley activities.

The focus of technology scouting was in two areas: (a) based on a deep understanding of existing products and services, we scouted for technologies that could result in significant improvement of those, and (b) scout for technologies that are adjacent to the core business and could create new business opportunities. The latter also has the potential to enter new areas that could eventually replace the existing “cash cows”.

The benefit of equity investments via venture capital are to enhance the relationship with a startup and its ecosystem, to improve the startup’s probability of success, while at the same time leaving the startup to develop naturally without suffocation by the large corporate entity. The equity investment is structured in way that gives certain preferential rights that can be exercised down the road, such as the right of first refusal, and others.

Silicon Valley is a microcosm of bigger changes that will occur later. Having first hand insight from research topics at local universities and research institutions, startup activities, venture capital flows, and industry conferences and gatherings allowed us to build an understanding of where things are going. This perspective was delivered to C-level company executives and was instrumental in shaping the strategic direction of the corporation.

Specifically, the Silicon Valley office recommended concrete strategies for integrating WiFi into a mobile service, avoiding WiMax and Fast-tracking LTE, embracing mobile advertising and accelerating entry into that business through VC investments, etc. These are just a small number of specific deliverables that had a major strategic and business benefit to the company.

Case Study 2: Contributions to the development of a tech innovation eco-system

In this case, the project included working with a country based innovation eco-system, including their venture capital arms, government R&D arms and academic and technology innovation institutions, to bridge them into the Silicon Valley eco-system, and define the best strategy for short, mid and long term cross-fertilization.

Taking into account this country aim to evolve towards becoming a Tier 1 R&D innovation technology hub, there was an opportunity to contribute to this evolution, strategically and tactically, based on the past experiences we have in the Silicon Valley innovation hub. Specifically, the goal was to assist with the development of a technology incubation initiative, from ideas inception to an early stage go to market strategy of select R&D initiatives. This did include:

- Work with select set entrepreneurs and soon to launch technology startups academics on a way to optimally jumpstart their ideas-to-startup process, and jointly build a path towards leveraging innovation towards commercialization success
- Work with them to analyze and review select R&D proposals, with market innovation in mind, based on the various goals and target markets
- Work with the potential early stage incubation fund and/or investors on the due diligence side.
- Work with the various actors of the overall innovation eco-system on putting the right elements in place for success
- Bridge in Silicon Valley venture capital process as part of the value chain engagement
- Develop relationships and explore synergies with the Silicon Valley ecosystem

The following approach has been taken to achieve this

- Spend some significant amount of time with the various stakeholders, working with them on building a path towards developing early stage ideas, with a path towards a startup venture.
- Develop a good remote working model to achieve optimal collaboration with the various stakeholders, with a focus on mentoring some of the new ventures, working with the venture capitalists to bring in high value funding, and bridging in into some key technology players that would be potential partners, channels to markets or in some cases, acquirers of such ventures
- Experiment with the above and refine based on progress, as far as model and engagement, tactically and strategically.
- The final outcome included very positive developments on various fronts, as stated by the initial objectives. Specifically, it did open other synergetic relationships with the Silicon Valley for various stakeholders, including the entrepreneurs, venture capitalists and government funding groups.

Case Study 3: Building a bridge to Silicon Valley

In this case, the partner was a hybrid public/private supported entity, that wanted to enhance technology innovation in their geography and act as a nucleus for entrepreneurship and startups. This included the layout of an innovation stimulus Fund. The stated goals included:

- Fast track plug into the Ecosystem: Access to ecosystem of startups, other Silicon Valley incubators, accelerators, angels, and solutions developers.
- Scout for high end technical and business development resources: Identifying and enrolling subject matter experts in specialized field areas, including subject matter experts and network of what would potentially become a Global Entrepreneurs in Residence (G-EIR)
- Designing Enablers: Custom pitch events and strategy consulting/trends. Answering the partner's specific needs around identifying and engaging with startups.
- Design of experimentation Sandboxes: A sandbox for innovation enablement based on the specific needs of the partner for technology experimentation (e.g. Fintech for regulators)
- Access to talent pool of knowledge worker in emerging technologies, on a global scale, with a strong Silicon Valley tie up

In addition to the local in-country structure, we added an integral "bridge to Silicon Valley" plan. This constituted an office in Silicon Valley operated by Xona Partners under the brand of the partner, with models of reverse white labeling over time.

The functions performed in the Silicon Valley Office included:

- Help review startups that apply to the in-country Program by providing a Silicon Valley perspective and evaluation methodology.
- Assess the technology in the proposal and any prior art in Silicon Valley and if needed advise on modification to the business plan.
- Assess market traction potential of the proposal, any possible partnerships, and if needed advise on modification to the business plan.
- Help with the investment recommendation to the Fund.
- Suggest possible additional co-investors from US ecosystem.
- Source guest speakers from Silicon Valley and the US to visit the Program HQ.
- Host visitors from Program HQ and organize educational tours of Silicon Valley for Startups and executives from the Program, the Fund, and other related executives and government officials.
- Create and manage an entrepreneurship syllabus for startups in the Program.
- Advise Program startups on goals and objectives as they evolve.

- Assist with business opportunities in the US.
- Assist with follow-on funding from the US and the Fund.
- Help Program startups open office or virtual office in Silicon Valley.
- Help Program startups incorporate in the US when that is beneficial.
- Assist in M&A when opportunity arises.
- Use extensive connections with Silicon Valley startups and established companies to promote using in-country resources and the startups supported by the Fund to partner with and to do trials in country or open branches in country.

Conclusion

Silicon Valley is here to stay for the foreseeable future. It will remain a constant draw for inspiration to a continuous flow of innovators from around the globe. Yet, for the corporations, governments as well as incubators and accelerators wanting to have a valley presence, this is not easy task. Finding the right model is fundamentally important, prior to diving into execution. We, as a Xona Partners team, aim at easing this process and having it tailored based on the specific innovation needs. This paper described the rationale for such approach. Real world illustrations were provided, highlighting successful recent engagements, where the end goal was to not only set foot in the valley, but do so in a way that ended up providing the most optimal results over the shortest amount of time, and in the most economical manner, as highlighted from the experiences of our partners.

XONNA PARTNERS

A Perspective on Multi-Access Edge Computing

Frank Rayal

January 2017

Overview

The convergence of the Internet and telecommunication networks is igniting a debate on how to cost effectively meet the performance requirements of different services. Edge Computing, which places compute and storage resources at the edge of the network, is a technology at the heart of this debate. It promises to bring many benefits to end users, but its implementation in mobile networks has to overcome a number of challenges.

The interest in Edge Computing is enforced by the emergence of 5G wireless access technology with applications in varied vertical markets such as automotive, health, energy, education and many others. This widens the addressable market for service providers beyond the traditional consumer-centric business model. This is a fundamental shift that affects many aspects of business operations.

In our analysis of Multi-access Edge Compute (MEC) – the implementation of Edge Computing in wireless networks – we conclude that mobile network operators (MNOs) are still debating which approach to take. Some view MEC as integral part of 5G networks where it would be coupled with a new network architecture. Others view MEC as a tactical expediency in certain applications. A strategic view of MEC and its role in the network is missing to date, largely due to the multitude of applications and beneficiaries.

Equipment vendors are making provisions for MEC in the design and architecture of their solutions. Vendors have embraced a flexible network architecture that centralizes specific functions of the radio access network to improve performance of heterogeneous networks. This is coupled with steady progress in network virtualization that will facilitate the implementation of MEC. Nevertheless, vendors have a major challenge in getting the value proposition of MEC fully exposed, due to the existence of a large number of use cases and stakeholders of MEC.

The applications for MEC correlate closely with vertical markets, which have different service requirements. As a result, the implementation of MEC could accelerate with market adoption of these applications and technologies: IoT connectivity, small cells, new spectrum regimes, and technologies such as virtualization and network slicing. Such technologies raise the opportunity for third parties to deploy MEC in private networks, independent of the mobile network operators. Hence, MEC becomes an integral part of a service offering that differentiates from that provided by MNOs. Therein lies an opportunity for new entrants to leverage flexible business models.

Key Takeaways

1. Edge Computing is necessary to meet the requirements of 5G applications and allows service providers to address the needs of vertical markets.
2. Implementation of Multi-access Edge Computing is coupled with a compelling business case that is absent today as service providers develop their strategy on how to best address vertical markets.
3. Equipment vendors must remain flexible on how to implement MEC in order to meet a range of potential applications with varying requirements and market potential.

MEC Definition

MEC moves compute and storage functions closer to the end user at the edge of the network and away from the core (Figure 1). This improves the response time of applications (reduces latency) and reduces the amount of data traversing the transport network between the core and the MEC server location. Distributing compute and storage results in additional cost to the service provider. Consequently, the location and sizing of compute and storage elements are key design factors that need to be carefully balanced.

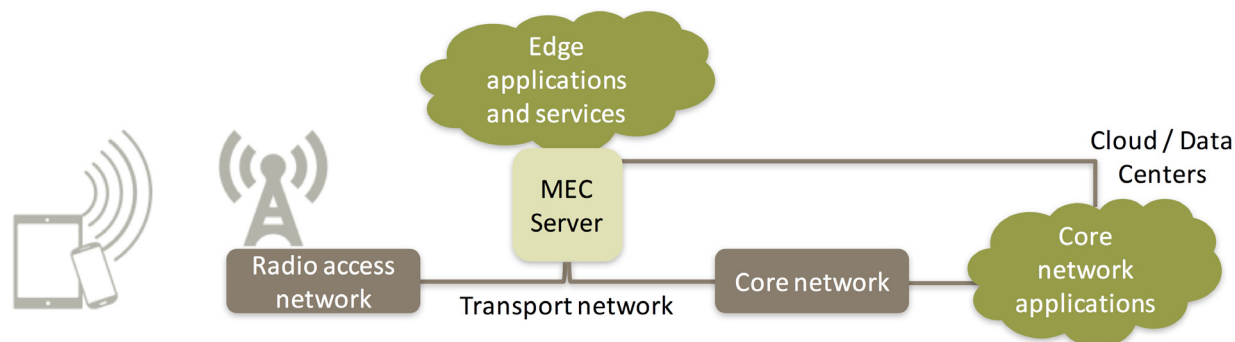


Figure 1: MEC architecture overview.

MEC impacts the architecture of the mobile network, which centralizes important functions such as billing and legal intercept. Having compute and storage capability at the edge entails a non-trivial expansion of these functions to the network edge, particularly within the framework of today's highly distributed LTE networks.

The Applications

Applications that benefit from MEC include those with one or more of the following requirements:

- High responsiveness, low latency and near real-time operation
- Data caching
- Context-aware services
- Location-aware services
- Heavy computation applications
- Data transformation and transcoding
- Extended battery operation

This includes the following applications:

- Enterprise applications including asset tracking, video surveillance and analytics, local voice and data routing.
- Augmented and virtual reality.

- Multimedia content delivery where video can specifically benefit from caching and transcoding.
- Retail services including ad delivery and footprint analysis in shopping malls among other applications.
- IoT applications which can be divided into two categories:
 - a) Massive IoT connectivity where MEC streamlines device connectivity with the core network to reduce overhead communications and improves response time.
 - b) High-responsiveness applications where low latency is critical. This includes smart grid switching of power and alternative energy supplies, and fault detection applications.
- Critical communications: this category includes multiple applications in various sectors
 - a) Traffic safety and control systems.
 - b) Precision farming using autonomous vehicles and real-time analytics.
 - c) Industrial IoT applications for monitoring and time-critical process control.
 - d) Automotive applications related to hazard warning and cooperative autonomous driving.
 - e) Healthcare applications requiring high responsiveness.

Many of the above applications can only be implemented with MEC, the only way to provide sub 1 msec compute to the network edge.

MEC Drivers and Dependencies

MEC integrates with a number of technologies leading to a scenario of mutually enforcing adoption. Thus, the greater traction these technologies attain in the market, the more relevant MEC would be over and beyond its baseline. The following technologies help define the future for MEC:

- 1) **Network virtualization:** network function virtualization (NFV) and software defined networking (SDN) are two key technologies whose implementation significantly reduces the barriers to entry for MEC. Specifically, the application of NFV in the radio access network (RAN) is important. The major vendors have embarked on a process of redefining their RAN solution architecture to incorporate NFV and to readily provide a platform for MEC implementation.
- 2) **Small cells and heterogeneous networks:** MEC allows customized services in various use cases such as enterprise and venue applications (e.g. shopping malls, stadiums, and airports). The emergence of shared spectrum regimes such as the 3.5 GHz Citizen Broadband Radio Service (CBRS) provides a market opportunity for small cell networks to ramp up. Similar scenario can be expected with unlicensed band LTE technology (MuLTEfire). Small cell networks in shared and unlicensed spectrum need not be deployed by the MNOs, but can be deployed by private enterprises thus creating an opportunity to offer differentiating services.

- 3) **IoT connectivity:** Applications in the Industrial Internet is a major potential driver for MEC as it allows support for lower cost devices that packs less processing than otherwise required (i.e. thin devices). This results in lower latency and faster response. Applications in Industrial Internet have specific requirements related to latency, location, processing, etc. that MEC could fulfil effectively.
- 4) **Network Slicing:** This is a 5G technology that relates to provisioning instances or personalities of the network to serve applications with specific performance criteria. Network slicing leverages network virtualization concepts to create or remove network slices based on demand. While the full implementation of this technology is still a few years away, it integrates well with MEC where both technologies contribute to meeting the quality of service and experience subscribed to by the user.

The MEC Ecosystem & Business Case

Unlike other technologies, MEC opens up the possibility to change the telecom value chain by inserting new players including the MEC service provider and application developers (Figure 2). It also offers the potential to change how content providers and OTT players deliver their services. MEC allows service providers to capitalize on new business opportunities, such as applications catering to vertical market requirements, which leads to new business dynamics among players in the value chain.

For instance, MEC allows OTT and content providers to offer better service to end users, but the cost of the MEC infrastructure is borne by the service provider. How the future relationship between OTTs and content providers with telecom and MEC service providers will shape up is an open issue with multiple possible outcomes depending on the application.




Figure 2: MEC value chain.

The business case for MEC has high variance where, in addition to the wide range of possibilities on the revenue side, there is a wide variance in cost. The major factor for cost for MEC is the proximity of MEC servers to the network edge. More servers located at the network edge will result in increased performance, and cost. On one extreme, MEC servers can be placed at every base station. But this leads to the highest cost of deployment while leaving the number of users benefiting from MEC limited to those served by that base station. To address this, 5G networks are being architected to support multiple hierarchies whereby the MEC servers can be placed at an aggregation point between the core and the base station. Such an architecture captures the benefits from central offices that some service providers have, and is the central premise behind project such as Central

Office Re-architected as Data Center (CORD) and Mobile-CORD (M-CORD). Alternatively, it is possible to leverage the small cell gateway or controllers in heterogeneous network deployments (Figure 3).




Figure 3: MEC deployment scenarios.

Challenges to MEC Implementation

There are a few commercial and technical challenges to implementing MEC. Of the technology challenges, we note specifically:

- 1) Technical compatibility with the current network architecture. For instance, functions such as billing and legal intercept are located in the core network. However, MEC fractures that architecture as data flow does concentrates at the edge and does not reach the core. The question is then how existing networks would be re-architected to leverage the benefits of MEC?
- 2) Ensuring security and network integrity in order to provides an open environment for third party application developers to run services on the telecom service provider infrastructure.
- 3) Maintaining service over a number of radio access technologies that characterizes heterogeneous networks, such as LTE, Wi-Fi and future 5G technologies.

As for commercial challenges, the issues today concentrate on highlighting the business case for an open MEC environment to both service providers and potential beneficiaries of MEC such as enterprises. This cannot be done in isolation of the application on hand and is specific to different vertical players which makes the market evolution of MEC very selective. Another issue relates to the handling of content including digital rights and content access management, and encryption and storage of the content within the network.

Conclusions

Edge Computing is a necessary architecture to meet 5G requirements, and enables service providers to enter vertical markets. This makes Edge Computing a cornerstone architecture for any service provider with plans to serve vertical markets. MEC, which represent the implementation of Edge Computing in wireless networks, is an evolving architecture that benefits existing 4G networks as well. While equipment vendors develop solutions that accommodate Edge Computing, service providers remain undecided on their approach to MEC. A chief reason for this is the absence of validated applications and a compelling value proposition for prospective customers. Virtualization of the wireless networks will positively impact the implementation of MEC as it reduces the barrier to entry. Moreover, the advent of applications such as private networks will play a positive role in accelerating the adoption of MEC.

XO PARTNERS

Will Open Source Disrupt the Telecom Value Chain?

Frank Royal, Riad Hartani, Rolf Lumpe

September 2016

Overview

The rapid emergence of open source models in the design of telecom and network infrastructure systems is a trend that could drastically change the industry value chain and underlying competitive dynamics. Several strategic initiatives initiated by Cloud and Internet players, telecom operators, and disruptive startups portend a new paradigm reshaping this sector over the next few years. We foresee that the telecom infrastructure value chain will experience significant change in dynamics within the next 3 to 5 years, where leading players must master technologies related to Cloud development and deployment models. The vendor landscape will also change significantly, led by Cloud centric players. Service providers will also face significant business and competitive pressures, due to the strategic push of new business models by cash-rich Internet and Cloud players.

In this paper, we summarize our view on the potential of open source technologies and impact on business models in the telecom ecosystem. These views are based on Xona Partners' involvement in the development of telecom networks and Cloud infrastructure, which we intersect with the observations of leaders in these ecosystems.

Introduction

The design of networking equipment for telecom operators has always been the realm of specialized vendors, whose solutions were based on proprietary, in-house implementations of standards-based technologies. Attempts to open source some solutions in the 1990s and 2000s, e.g. routing operating systems, business operating systems and security systems had a timid effect on the overall industry, with a limited impact on the ecosystem. In the 2010s, three major trends appeared that brought in a new perspective to open source. First is the emergence of Network Function Virtualization (NFV) in the data center compute and storage environment, and a gradual evolution into the data center networking infrastructure. Second is the involvement of the large-scale Cloud and

Internet providers, who have for a long time designed most of their data center hardware and software in-house for competitive differentiation. The third trend is the fast evolution of open source in the areas adjacent to networking. This includes technologies such as: OpenStack and Docker's Cloud management and micro-service architectures; Hadoop and Cassandra's Big Data; and Jenkins and Spinnaker's DevOps continuous integration and delivery.

The Approach

To analyze the impact of open source on telecom, we divided the current and prospective telecom industry value chain into nine categories which include SPs, TEMs/SIs, semiconductor/baseband and software stack vendors, Internet, Open Source, and IT/Data center players, ODMs, and startup companies. We analyzed the impact of open source on each category supported by interviews with a representative sample in each category. We also analyzed different representative cases in the IT ecosystem to draw parallels with the telecom ecosystem.

The trends in open source intersect with shifting business dynamics that require telecom operators to adopt agile and service aware networks. Furthermore, the emergence of alternative wireless technologies is enabling new competing service providers including the insertion of the Internet and Cloud players in the telecom infrastructure value chain. These developments increased the momentum of open source telecom equipment solutions with the objective of increased agility, reduced time and cost of development and lower cost of deployment.

Open Source in Brief

Open source refers to the ability to access and modify source code, develop derived works, and sell or distribute software; i.e. open source does not imply free of charge. The construct of open source leads to collaborative communities, and consequently a philosophy in product development that is characterized by a relatively fast iterative process, where activities such as functional and interoperability testing are part of the development process. This contrasts the development process followed in telecom networks for standards-based equipment, which is characterized by a sequential 'waterfall' process that is well defined, but is relatively slow. Here, we like to note that while open source refers traditionally to software, it can apply to hardware as well in which case a reference design is shared in an open community.

Open Source Business Models.

Monetizing open source solutions can take different forms including the following common models:

- Offer complementary services or products to open source products such as support, maintenance, consulting, or hosting.
- Provide a commercial version or extension of open source products.
- Provide dual-licensing of proprietary solutions where a company offers its own proprietary software for use under either of an open source license or a paid commercial license.

The Present Landscape.

Hardware, or appliance, solutions make up the vast majority of telecom network infrastructure, a testimony of legacy services based on vendor-specific solutions. In recent years NFV, and to some extent Software Defined Networks (SDN), solutions began to appear in networks, starting with the outer perimeters in the OSS/BSS and services such as virtual operator enablement and IoT connectivity. NFV and SDN applications were then introduced into the core network when a few leading MNOs began implementing virtual EPCs and IMSs for commercial services. Now in its very early stages, this trend is expected to evolve and accelerate in the near future.

SDN/NFV solutions will see wider adoption and deployment as operators seek flexibility in developing and launching new services that are critical to their competitiveness, especially against the over-the-top (OTT) players. OTT players have leveraged IT infrastructure

virtualization and open source solutions to achieve economies of scale, cost efficiency, and service agility exceeding the established telecom service providers.

The Advent of Open Source in Telecom.

While virtualization provides a leap in flexibility over hardware-based networks, being transformative to business models and operations for both MNOs and vendors, virtualization solutions remain proprietary implementations that are optimized for performance. Open source solutions that build on SDN/NFV promise to open up the network to third parties, adding vitality to a mature market and stimulating innovation. A few open source projects were recently launched in telecom networks, such as the Carrier Open Compute Project (Carrier OCP) in January 2016 by AT&T, Deutsche Telekom, EE, SK Telecom, and Verizon. Carrier OCP builds on the OCP framework for data centers and extends the scope to the telecom infrastructure under the Telecom Infrastructure Project (TIP). TIP has the goal of bringing open source design models to hardware and software solutions that meet the requirements of telecom service providers. While this is still an early stage, it serves to highlight the efforts and momentum behind such initiatives.

Another example of a service provider led open source initiative is M-CORD, a joint project between the ON.Lab and The Linux Foundation, driven primarily by AT&T, SK Telecom, Verizon and NTT. In parallel to these initiatives, MNOs have transitioned certain aspects of their networks to open source. AT&T's ECOMP is one example which is related to the control, management and policy of the network. Another example is Open Source MANO (OSM) to which Telefonica made major contributions. There are so many open source projects today, that it is a challenge to assess which to participate in, contribute to, and more importantly which ones to develop solutions around. Moreover, it is important to note two fundamental aspects. The first is that open source has extended its reach from software-only to now include hardware in all its variants. The second is the heavy involvement of the Internet and Cloud players, such as Facebook with OCP and Google with M-CORD, to accelerate development and adoption of these technologies. The involvement of the Internet and Cloud giants in access technologies is a response to investor pressure on these highly-profitable, cash-rich companies for continued revenue growth – driving them to reach into lower, more cost-sensitive segments of the consumer market. Open source is therefore a vehicle to enable the development of applications and services across different market segments that otherwise would not be possible to achieve.

Motivations for Open Source in Telecom.

Open source projects are largely MNO-led initiatives with strong support from the Internet and Cloud players. The main reasons in priorities cited include:

- a) **Reduce vendor lock:** Consolidation of Telecom Equipment Manufactures (TEMs) has led to a few companies, such as Ericsson and Huawei, with overwhelming infrastructure market share. This impacts the innovation cycle and it becomes imperative for MNOs to stimulate innovation and creative solutions through open source.

- b) Morph cost models from capex to opex: The question of cost is complex, as MNOs are not necessarily expecting major reduction in the total cost of ownership from SDN/NFV-based solutions. What is certain is that in open source, as is the case with SDN/NFV-based solutions, the cost model is opex-based, which provides higher capital efficiency and is more responsive to network scalability, especially for new services such as machine connectivity.
- c) Enabling new services: The leading MNOs feel highly constrained within the confines of the existing network infrastructure. They seek the ability to deploy new services and features more cost effectively to improve their competitive positioning, especially against OTT services. The type of services MNOs seek vary according to region and range, from highly advanced applications such as V2X to more common ones such as rural connectivity.
- d) Stimulate and accelerate the innovation cycle: The leading MNOs participate heavily in standard activities to drive their vision into the process and ensure that the standard will meet their requirements. The 'waterfall' process is slow in the context of rapid technological innovation. Many standards exist, of which only a few are used. Open source as an iterative process is a means to accelerate the technology development and the deployment cycle.

The view among MNOs on open source is not universal, and there is divergence among leading Tier 1 service providers and others, who are more willing to take a wait-and-see approach.

Ecosystem Positioning on Open Source in Telecom.

TEMs who are a key part of the ecosystems, and often take on the system integration function, are largely ambivalent about open source projects at the current time; mainly because of uncertain financial benefits and large commitments. TEMs have invested heavily into product development, including optimization of complex interconnected sub-systems. They would argue that reliability, security and performance are paramount. Additionally, intellectual property rights form a significant source of revenue that TEMs will want to protect. They are currently evaluating potential loss/benefit scenarios for the transition to open source models. System integration, which is a critical function, would still be required irrespective of the approach to product development and deployment. Hence, open source can bring about a transformation in the telecom value chain that would result in a new division of functions. To kick-start the process, the MNOs themselves would have to lead the transformation, which is a challenging endeavor. New system integration entrants would need to have the financial and logistical strength to change the market, which is possible and more likely when a new application receives wide market traction to stimulate the open source model.

The Impact of Open Source in Telecom.

One fundamental trend is obvious - the gradual introduction of NFV/SDN solutions in telecom networks. With that, a gradual increase of open source components to build and deploy virtualized solutions. The consequence of this evolution is to morph the value chain

by raising system integration to the forefront, where different players will be positioned to build solutions around open source and provide end-to-end integration and deployment solutions. Although TEMs are best positioned to capture this activity in the early stages, the main threat comes from the players who possess full control over the virtualization, Cloud, and DevOps value chain that will form the cornerstone of the telecom services offering. The ability of these players to impact the market is stimulated by applications where virtualization is a cornerstone technology required to ensure cost effective operation. Consider, for example, IoT connectivity in wireless networks where core network elements are virtualized for scalability and cost efficiency. Other applications include enterprise services and small cell networks, particularly those operating in shared or license-exempt spectrum. The interest in open source is evolving in parallel with developments in 5G technology. 5G requirements and diversity of applications mandate a heterogeneous network where virtualization technology is a prerequisite to enabling concepts such as network slicing.

In our view, the telecom infrastructure value chain will experience a significant change in dynamics within the next 3 to 5 years, where leading players will have to master the technologies that are seen as adjacent to telecom today. Specifically, those that relate to deploying over Cloud infrastructure, agile application development, and efficient large data set management. All are areas where open source already plays a large role, which will extend to reach into telecom infrastructure. There are opportunities and challenges that will inject new vitality and innovative spirit into a market that is considered to be consolidated and mature.

Key Conclusions.

- Open source incursion in the telecom value chain is driven by telecom service providers and heavily supported by the Internet and Cloud giants.
- The main objective of open source is to provide service providers with a higher level of control over the network, and a flexible environment to quickly develop and launch services to generate new revenues.
- Cost is a secondary consideration for service providers, while it is a foremost consideration for the Internet and Cloud giants, who seek to lower the cost of Internet access to increase market penetration.
- The telecom infrastructure value chain is set to experience significant changes in dynamics within the next 3 to 5 years, where leading players have to master development and deployment technologies related to Cloud, Data and DevOps models.
- The vendor landscape is likely to change significantly, with the Cloud centric players likely to lead.
- The service providers will face significant business and competitive pressure due to the strategic push of new business models by cash-rich Internet and Cloud players.

Acronyms

BSS	Business Support System
ECOMP	Enhanced Control, Orchestration, Management & Policy
IoT	Internet of Things
LTE	Long Term Evolution
M-CORD	Mobile Central Office Re-architected as a Datacenter
MANO	Management and Orchestration
MNO	Mobile Network Operator
NFV	Network Function Virtualization
OCP	Open Compute Project
ODM	Original Design Manufacturer
OSM	Open Source MANO
OSS	Operations Support System
OTT	Over-The-Top
SDN	Software Defined Networks
SI	System Integrator
SP	Service Provider
TEM	Telecom Equipment Manufacturer
TIP	Telecom Infra Project
V2X	Vehicle to Anything

XOONA PARTNERS

RAN Virtualization: Unleashing Opportunities for Market Disruption

Frank Rayal

June 2016

Table of Contents

List of Figures	3
List of Tables	3
Overview	4
The Genesis	4
A Disruptive Idea	5
The Challenges	6
The Solutions	6
Categorization of Architectures	8
Market Trends	11
Vendors' Strategies	11
Telecom and Internet Ecosystem Convergence	12
Impact on DAS and Small Cell Ecosystems	12
Fixed Access Service Providers and Neutral Hosts	12
Evolution Towards 5G	13
The Financial Business Case	13
The Ecosystem	13
Conclusions	14
Acronyms	15

List of Figures

FIGURE 1	VIRTUAL RAN: BASEBAND VIRTUALIZATION.	5
FIGURE 2	FUNCTIONAL PARTITIONING OF THE LTE PROTOCOL STACK.	7
FIGURE 3	RAN ARCHITECTURE DEFINITIONS.	9
FIGURE 4	SIMPLIFIED VRAN ARCHITECTURE.	10
FIGURE 5	FUNCTIONAL SPLIT TRENDS FOR LTE.	12
FIGURE 6	CLOUD RAN ECOSYSTEM.	14

List of Tables

TABLE 1	BACKHAUL AND FRONTHAUL REQUIREMENTS FOR A 20-MHZ 2X2 MIMO LTE CARRIER.	6
TABLE 2	OVERVIEW OF FUNCTIONAL SPLIT CHARACTERISTICS.	8
TABLE 3	RAN ARCHITECTURE DEFINITIONS.	11

Overview

Radio Access Network (RAN) virtualization is a highly disruptive technology that will radically impact how wireless services are delivered. It will change the current ecosystem and market structure; altering the way MNOs plan and roll out new services by providing a scalable, on demand alternative to the traditional architecture. Dedicated, on-site hardware to power the RAN is becoming expensive to build-out and maintain especially as more cell sites are required to keep up with capacity demand. Virtual Radio Access Networks (vRAN) moves the baseband modules away from the radio at the cell site to a data center. This enables intelligent scaling of computing resources as demand on capacity fluctuates, while reducing site lease costs, energy usage, and maintenance expenses. The evolution of LTE and advent of 5G networks increases bandwidth requirements further. This makes increased fronthaul requirements and the inflexibility of the legacy CPRI serial interface the primary challenges to vRAN deployments. Resolving the fronthaul challenge enables the Internet giants and fixed access service providers to enter the wireless market with lower cost basis, a move that is highly disruptive in a market dominated by telecom incumbents entrenched through massive equipment install-base.

The Genesis

Mobile network operators (MNOs) in Japan and Korea were first to centralize the radio access network by moving base stations baseband units to fiber centers, leaving only the remote radios and antennas at the cell site. This network architecture is possible provided fiber is available to link the baseband units to the remote radio – a link called fronthaul. Operational cost savings from this architecture range between 30 – 40% due to lower site lease, simplified support and maintenance, as well as lowered energy expenses. Operators without their own fiber assets would find it cost prohibitive to implement this architecture because of the high fronthaul performance requirements of legacy protocols used to connect the baseband to the radio (e.g. CPRI). Improvements to this link will make fronthaul feasible to service providers without their own fiber assets.

At the turn of the decade, LTE deployments were burgeoning and data traffic was doubling year over year. Unfortunately for MNOs, the average revenue per user (ARPU) did not increase, falling in many markets and leading to lower EBITDA margins. Some of MNOs, such as China Mobile, saw virtualized RAN as an opportunity to lower costs and improve financial performance. Together with other Asian operators, China Mobile promoted the concept of Cloud RAN, which virtualizes the centralized baseband processing to achieve further cost savings. The term Cloud RAN has since become a buzzword, and many vendors with different solutions began using the term liberally, a few with little relationship to actual Cloud RAN. We will use the term vRAN to denote a fully centralized and virtualized baseband implementation (Figure 1).

On top of cost savings, vRAN also brings performance benefits. This is owing to features such as coordinated multipoint and network MIMO, which become possible due to centralization, and are utilized to lower interference and improve throughput. The result is enhanced user experience, especially at the cell edge where performance is most lacking (up to 100% throughput gain at the cell edge has been demonstrated in field trials). In fact,

centralization becomes more important in heterogeneous networks (HetNets) where low-power small cells are deployed in the service area of high-power macrocells. Centralization reverses the LTE distributed architecture which places the entire protocol stack at the base station leading to high overhead and timing requirements for coordination among base stations to mitigate interference. Future network architectures planned for 5G intend on implementing a flexible architecture, where part of the intelligence is centralized to reduce the coordination overhead.

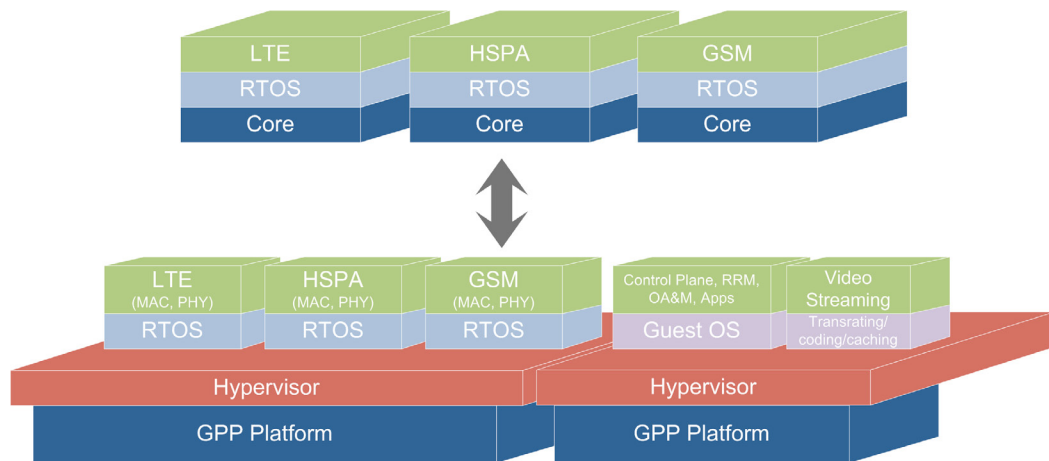


Figure 1 Virtual RAN: baseband virtualization.

The gain associated with virtualization is based on leveraging the cost structure and economies of scale of the IT/data center industry. Furthermore, the scalable and elastic properties of virtualization allow deploying processing power to provide capacity on demand when and where it is required in sharp contrast to distributed hardware architecture that is designed for peak capacity.

A Disruptive Idea

Virtualization decouples the software from hardware, enabling the use of commercial servers in the network. This profoundly alters the way MNOs plan, design, procure and roll out new services. They would no longer need to purchase hardware-optimized base stations from specific telecom equipment manufacturers (TEMs). Instead they would only need software and general purpose servers in data centers to run the wireless protocol stack as an application to power any remote radios on demand. Other applications can run on the same infrastructure to provide value added services, such as video optimization, caching and localization. TEMs could provide their applications in a software as a service (SaaS) setting, with an OPEX-based pricing model, instead of the CAPEX-dominant model of today. MNOs could control and manage large networks more efficiently to enable a HetNet architecture. Because wireless capacity is not in demand at peak level at all locations at the same time, MNOs could save substantial expenses by multiplexing wireless capacity to increase operational efficiency and reduce capital costs. The RAN market structure will be radically changed, altering the balance of power between vendors and operators; leading new entrants into a market that's becoming highly consolidated. Such is the disruptive nature of virtualization in the RAN.

The Challenges

The major challenge to implementing vRAN is the fronthaul interface between the baseband units and the remote radio. CPRI is the most common interface, which was designed in 2002 before the centralized architecture was advanced. It requires 10x the capacity of an LTE backhaul channel, which makes it prohibitively expensive for operators who don't own fiber assets. Unlike backhaul, CPRI fronthaul cannot be statistically multiplexed so its capacity requirements increase proportionally with the number of LTE carriers used. CPRI also has tight requirements for synchronization, latency and jitter that are difficult to meet when there is no direct connectivity between baseband and radio. As a result of these factors, fiber becomes the only media capable to implement fronthaul. While this is possible, especially as the cost and transmission capabilities of optical transceivers have been on a steep improvement curve, it remains a challenge to many operators who don't own fiber or where fiber penetration is thin.

Table 1 Backhaul and fronthaul requirements for a 20-MHz 2x2 MIMO LTE carrier.

# of Carriers	Backhaul (Mbps)	Fronthaul (Mbps)
1	236	2,547
3	248	7,641
6	496	15,282
9	744	22,923
12	992	30,564

A second challenge pertains to virtualization. The wireless protocol stack includes computationally intensive functions that are inefficient to run on general purpose processors (GPPs). Devices such as FPGAs, ASICs and SoCs are more efficient, and provide real-time response capability, which is required by some RAN functions. Such challenges are beginning to dissipate as new, more powerful, GPPs with vector acceleration functions are becoming available on the market. Additionally, there are different implementations of virtualization that can solve these challenges such as offloading complex functions to acceleration engines. It is now clear that challenges due to virtualization could be overcome as demonstrated in recent PoCs, where performance was near that of hardware-based implementations.

The Solutions

The solution to the fronthaul challenge takes different paths depending on the objective. If the goal is to ensure compatibility with installed base of remote radios, CPRI compression techniques may be used. These typically achieve between 50% – 66% savings in bandwidth. Alternatively, the protocol stack can be divided, with some functions virtualized at the center and others performed at the cell site. The functional split of the protocol stack trades off potential performance enhancement against fronthaul latency and capacity requirements (Figure 2, Table 2).

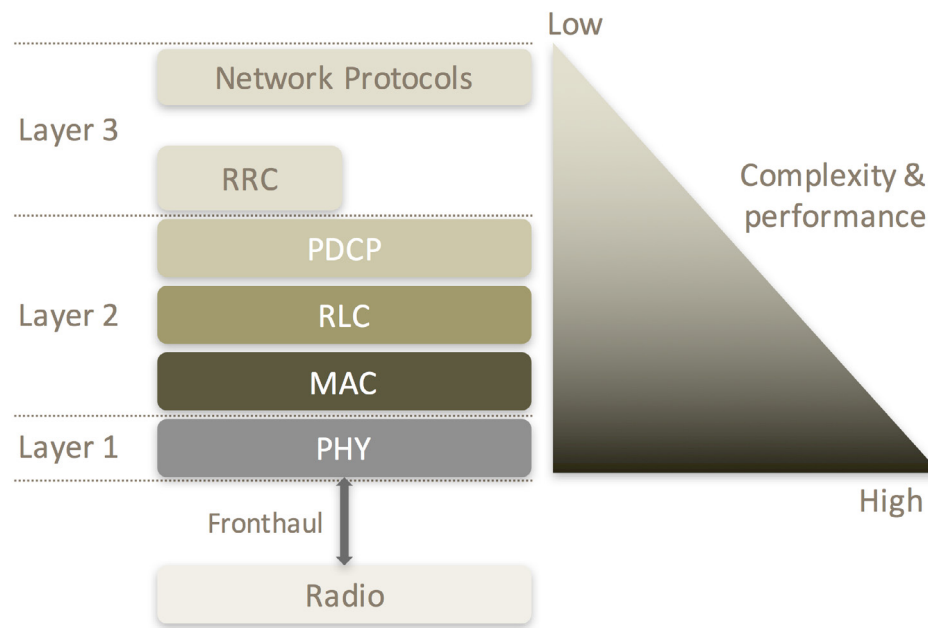


Figure 2 Functional partitioning of the LTE protocol stack.

While such approaches cater to accommodating legacy networks, it is possible to design new interfaces optimized to meet the requirements of future networks (high scalability, low cost). Such interfaces bring about the full benefits of RAN virtualization and revolutionize the wireless infrastructure market. While the technology has been demonstrated, achieving consensus in the industry is more challenging as incumbents work to protect their market share and position. Several industry forums have initiated studies to engineer a new interface – these efforts are still at a relatively early stage.

Table 2 Overview of functional split characteristics.

	High Functional Split	Low Functional Split
Fronthaul requirements	1 – 2x the capacity requirements of backhaul	Same as CPRI requirements, if CPRI is used
Performance enhancements	Limited in comparison to low functional split but better than a fully distributed architecture	Maximum performance enhancements though CoMP and network MIMO techniques
Cost of implementation	Low cost in comparison to distributed architecture	High cost if CPRI fronthaul is used
Compatibility with installed-base	High compatibility with current install-base of equipment: could be implemented with additional network elements	Limited compatibility with current install-base of equipment
Disruptive potential	None – similar fundamental building blocks to the current distributed architecture	Disruptive potential requires an efficient packet-based interface. Low disruptive potential with CPRI

Categorization of Architectures

In an effort to improve performance of the distributed LTE architecture in HetNets to meet future capacity demand, equipment vendors are beginning to centralize parts of the protocol stack. Virtualization is implemented in some centralized designs, but not all. This has led to a bifurcation of architectures that diluted the term Cloud RAN. From its original definition of fully centralized and virtualized air interface protocol stack, Cloud RAN is now even used to refer to solutions that include neither centralization nor virtualization. We introduce the following definitions while recognizing that different implementations exist within each category (Figure 3, Table 3):




Figure 3 RAN architecture definitions.

Virtual RAN: An architecture where general purpose processors and servers are used to run air interface protocol stack in a central location (Figure 4). Various architectures and implementations of vRAN exist:

- Architecture where all layers of the air interface protocol stack run on GPPs located in a central location.
- Architecture where non-real-time functions in Layer 2 and Layer 1 run on GPPs while real-time functions run on hardware accelerators.

Some implementations run the protocol stack on a processor without capabilities for pooling and load-sharing of resources (i.e. bare metal).




Figure 4 Simplified vRAN architecture.

Hybrid RAN: A split baseband architecture where some modem functions run on GPPs in the center while other baseband functions, such as Layer 1 or parts of Layer 2, run on programmable and hardware devices, such as FPGAs, DSPs, NPUs ASICs and SoCs, at the remote radio. The split can occur at different locations and is a vendor specific design. Hybrid RAN is an architecture that optimizes cost and performance but does not have the same disruptive potential as vRAN.

Clustered RAN: An architecture where baseband modules are located in a central location as is done in today's base station hotels. The air interface protocol stack runs on programmable and hardware devices. This is the most basic form of centralization, and is targeted for OPEX reduction in certain Asian markets. It is also used for practical considerations in other parts of the world where it is not possible to collocate the baseband with the remote radio due to different considerations such as space and access. Clustered RAN is the name given by SK Telecom to Phase 1 of their roadmap to implements vRAN.

Centralized RAN: An architecture where the baseband modules are located in a central location, similar to Clustered RAN, but with two variations:

- a. All the baseband functions of the air interface protocol stack are centralized (full centralization). In this case, the difference from Clustered RAN lies in the integration of baseband processing to save cost among different modems and to improve performance through coordination of resources.
- b. Part of the upper layers of the protocol stack are centralized while the lower layers are distributed at the remote radio (partial centralization) – essentially a split architecture without virtualized baseband.

In either case, the baseband processing is based on programmable devices running all air interface modem functions. The architecture supports a 1:1 relationship between a

radio and its baseband modem. GPPs may be used to run Layer 3 functions in addition to different applications.

Table 3 RAN architecture definitions.

		Architecture: Baseband Centralization		
		Centralized	Split	Distributed
Technology: Baseband Virtualization	Virtualized	Virtual RAN <ul style="list-style-type: none"> • Pioneered by startups • High potential for market disruption • Likely lead deployments in local-area coverage use cases (venues) 	Hybrid RAN <ul style="list-style-type: none"> • Supported by major vendors in wide-area deployments with a functional split high in the protocol stack 	Distributed RAN Architecture used in 490+ commercial LTE networks.
	Not Virtualized	Clustered RAN or Centralized RAN <ul style="list-style-type: none"> • Deployed on wide-scale by leading carriers in Korea and Japan for network OPEX savings • Deployed in select installation by operators worldwide for different reason: site acquisition challenges, zoning, security, power availability, theft prevention, etc. 		

Other terms are used in the industry to denote a level of coordination among base stations for interference management such as Cooperative, Collaborative and Elastic RAN (Ericsson) where the baseband processing is not necessarily virtualized. They can be classified according to one of the above categories.

Market Trends

Vendors' Strategies

Major equipment vendors are focusing on Hybrid RAN architectures that centralize and virtualize the upper layers of the protocol stack, typically the PDCP layer as it is a straight forward migration that utilizes existing infrastructure (Figure 5). This functional split allows the implementation of dual connectivity small cells, which improves mobility management in HetNet deployments.

Startup pioneers are leading in vRAN implementation, where different designs have emerged that promise to reshape the market landscape. vRAN lends itself to new ways of deploying small cells and distributed antenna systems (DAS).




Figure 5 Functional split trends for LTE.

Telecom and Internet Ecosystem Convergence

Behind the vRAN pioneers stand major Internet players such as Facebook, who initiated the Telecom Infrastructure Project (TIP) to explore the benefits of vRANs and its potential to reduce the cost of connectivity. TIP participants joined the Open Compute Platform (OCP) which is a 5-year old initiative on data center technologies for telecom companies. This points to the confluence of the Internet/compute world with the telecom world which has significant ramifications.

Impact on DAS and Small Cell Ecosystems

Deployment of vRAN is likely to be driven by venues and indoor applications, where demand for capacity is highest. This would precede deployments in macrocells, where there is already a large install base of LTE equipment in over 490 networks worldwide and a change in architecture is unlikely to occur before a major technology upgrade to 5G. The vRAN market will take off, provided the fronthaul connectivity requirements are similar to those of backhaul. vRAN would be a substitute for small cells and DAS, which is not optimized to support MIMO technologies, a leading feature in LTE (4x4 MIMO is a key feature of LTE-Advanced Pro; 3GPP Release 12 & 13). This development means greater overlap and interdependency between DAS vendors and TEMs.

Fixed Access Service Providers and Neutral Hosts

As LTE expands to unlicensed bands (e.g. 5 GHz) and shared spectrum bands (e.g. 3.5 GHz CBRS and 2.3 GHz), third parties will have the option to roll out LTE services there, concentrating on the indoor and venue markets. This allows companies with fixed assets such as fiber or cable, as well as neutral hosts, to enter the access service market with wireless solutions complementing those of the MNOs who own the wide-area coverage market.

Evolution Towards 5G

RAN Virtualization is a major topic as the definition of 5G networks emerges with varying use cases including extreme broadband, massive machine-type connectivity and ultra-reliable communications. The ability to run services at the network edge to optimize bandwidth utilization and user experience requires a configurable architecture. The scale which 5G networks are required to support can only be implemented cost effectively with a scalable and elastic network architecture. RAN virtualization provides this capability. However, as 5G incorporates millimeter wave bands for access services, different architectures will be in play as millimeter wave systems rely on large antenna arrays to achieve the desired coverage range.

The Financial Business Case

Analysis of different RAN architectures shows that the centralization of baseband leads to high operational cost savings in Asian markets (26%). This is due to the structure of cell site leases, limited availability of space at the cell site, and high energy costs. In North America, the structure of site leases is beginning to change. Energy costs are relatively low, such that the business case for vRAN would not be positive in all cases, especially as dark fiber will be required to meet the requirements of CPRI fronthaul. This results in high financial uncertainty and risk that deployment requirements can be met.

In HetNet deployments, fronthaul can overcome the advantage of wireless backhaul cost effectiveness (\$/Mbps) only if we consider high utilization of the remote cell. While Virtual and Hybrid RAN boost capacity, the average utilization of small cells over time is generally low, which erodes the return on investment. This issue is endemic to the HetNet architecture irrespective whether it is based on small cells or low-power remote radio.

In HetNets, fiber fronthaul is attractive in connecting remote small cells that are close to the macrocell. This is where interference between the HetNet layers is highest due to proximity. The breakeven point is about 75m: any remote cell at greater distance than 75m is better connected through wireless, if possible.

The major financial implications with vRAN is with regards to capital expenses. CAPEX reduction is driven by the baseband pooling gain of vRAN, however, that will depend on a number of factors. Primarily CAPEX savings depend on the deployment scenario and size of vRAN cluster, which is an MNO design option. Among other factors is the pricing model from vendors.

The Ecosystem

The Cloud RAN ecosystem comprises a wide cross section of vendors from the entire wireless ecosystem (Figure 6). However, we consider that a critical element of the ecosystem includes the Internet giants who are looking to reduce the cost of access to reach more subscribers and provide better quality of OTT services. Another important element of the ecosystem are the cable and fiber operators, whose fiber and other fixed access assets will have a major role in providing fronthaul services. These service providers already operate Wi-Fi as an extension to their fixed access services and some have looked

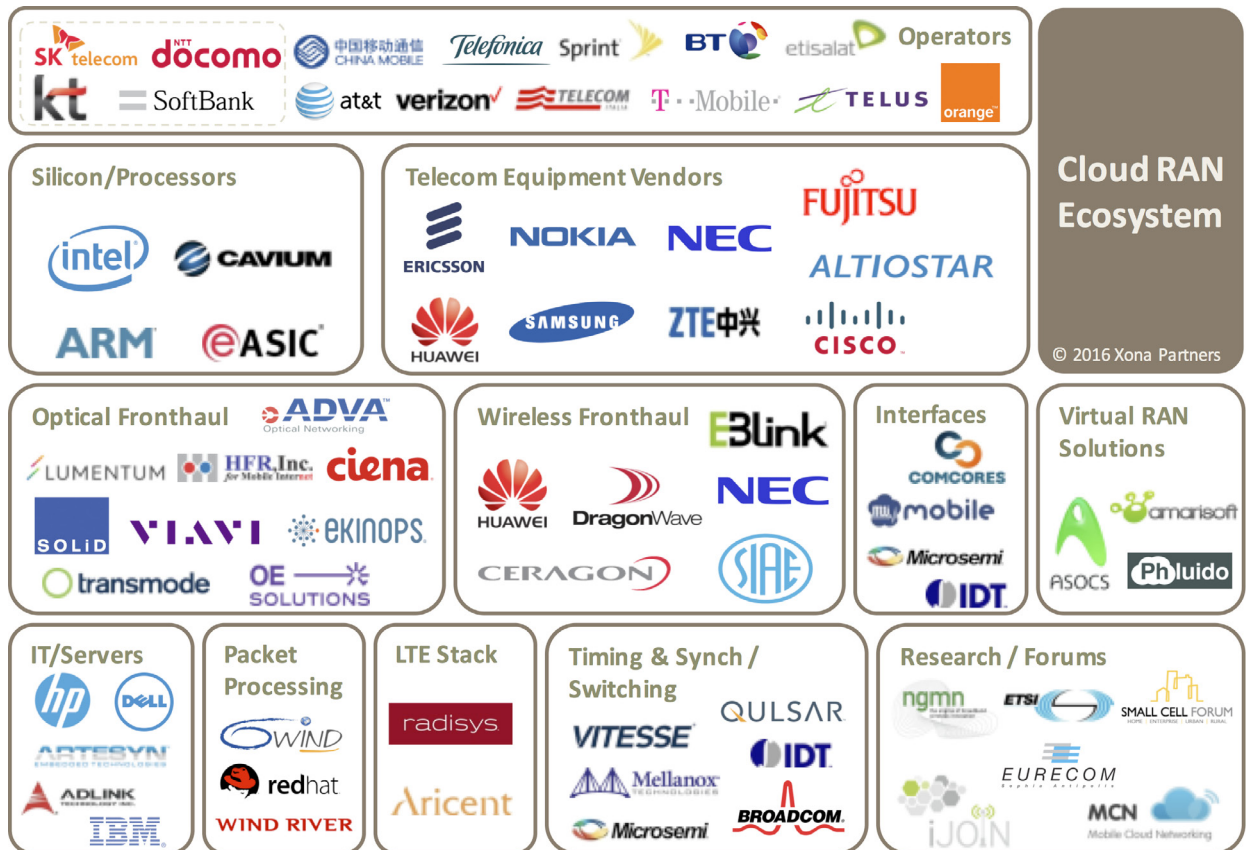


Figure 6 Cloud RAN ecosystem.

Conclusions

vRAN is a forward facing disruptive technology that is rapidly becoming more feasible as it garners support from Internet giants and startup pioneers. Current architectures being pursued by the TEMs, such as Hybrid RAN, will allow MNOs to improve the performance of HetNets specifically related to interference and mobility management, but will fall short of having a disruptive impact on the industry. Disruption will come from vRAN technologies when the fronthaul challenge is solved. This will alter the MNO-TEM relationship and market structure, and will allow new entrants into the market such as the fixed access service providers who can leverage their infrastructure for fronthaul services. The advent of RAN virtualization becomes especially potent when coupled with shared spectrum regulations, which increases the service possibilities and market opportunity.

Acronyms

3GPP	Third generation partnership project
5G	Fifth generation
ARPU	Average revenue per user
ASIC	Application-specific integrated circuit
CAPEX	Capital expenditure
CBRS	Citizen Band Radio Service
CoMP	Coordinated multipoint
CBRI	Common Public Radio Interface
DAS	Distributed antenna system
DSP	Digital signal processor
EBITDA	Earnings before interest tax depreciation and amortization
FPGA	Field programmable gate array
GPP	General purpose processor
HetNet	Heterogeneous network
LTE	Long Term Evolution
MIMO	Multiple input multiple output
MNO	Mobile network operator
MVNO	Mobile virtual network operator
NPU	Network processing unit
OCP	Open Compute Platform
OPEX	Operational expenditure
OTT	Over-the-Top
PDCP	Packet data convergence protocol
PoC	Proof of concept
RAN	Radio access network
RLC	Radio link control
RRC	Radio resource management
SoC	System on chip
TEM	Telecom equipment manufacturer
TIP	Telecom Infrastructure Project
vRAN	Virtual radio access network

XOONA PARTNERS

The State and Future of The Home Automation Market

Frank Rayal
Riad Hartani
Eric Chan

January 2016

Introduction

The home automation market is undergoing a progressive transformation propelled by the proliferation of smartphones and tablets. In addition to cellular technologies, home automation devices integrate different local and personal area technologies to connect among peripherals. This led to a new phase of evolution in home automation systems where wireless technologies enable connectivity for monitoring and control from anywhere at any time. Home automation solutions have broken through the early-adopter market phase. Mass market adoption on the other hand is yet to materialize leaving a great potential ahead for the next phase of development in a very dynamic market that's in the process of being defined. In this paper, we outline the main characteristics of the home automation market and expose trends are shaping the market, raising challenges, and creating new opportunities.

Market Characteristics

The home automation market comprises multiple segments, including:

- a. Lighting control (e.g. switches, dimmers)
- b. Security & access control (e.g. video surveillance, intrusion detection)
- c. HVAC control (e.g. thermostat)
- d. Entertainment control (e.g. home theater)
- e. Outdoor control (e.g. landscape)

We observe the following characteristics of the market for this market which outline the dynamics among various stakeholders:

Silo segments: The market is siloed into segments without support for unified interface or interaction. For example, HVAC is independent from security control systems leading to different user experience.

Fragmented use cases: Distinct use cases and applications result in fragmentation across multiple fault lines including markets and technologies.

Uneven adoption: Some market segments are more mature than others, in part due to market and distribution channels. For example, HVAC and security control are well established while lighting control is emergent propelled by regulatory requirements and incentives especially in EU countries.

Non-interoperable technologies: Vendors select the technology and protocol stack that best meet the application requirements leading to a proliferation of non-interoperable systems. Attempts are underway to bridge this gap through industry alliances and organizations which began taking shape in late 2013 and has accelerated since.

Security risks: Security shortcomings are alarming across a wide range of products. Examples include lack of enforcement of strong passwords, nonexistent support for

mutual authentication, or absence of protected accounts against brute-force attacks¹. Mobile applications are specifically vulnerable with estimated 20% not using encrypted communications to the cloud².

Inconsistent performance: Reliable connectivity is lacking due to a complex deployment scenario where signals could easily be blocked or be subject to interference from other systems operating in the same spectrum.

Market Players

Home automation is an active market with many players approaching it from different angles:

Technology giants: Apple (HomeKit, iOS), Google (Nest, Android), Samsung and Microsoft (Xbox, Windows) best exemplify this segment. These companies leverage the operating system of mobile devices, their incumbency in the Internet platform business, and the Cloud infrastructure to expand into the connected home market. New players are also entering this space with significant foreseen growth include Alibaba, Amazon and Xiaomi.

Industrial conglomerates: GE, Honeywell, Phillips, Schneider, and others, manufacture a wide range of appliances and home devices. Their centers on ensuring interoperability with home automation hub vendors as they seek to make their solutions as widely available to the market as possible. Some of these companies have decided to enter into the home hub market (e.g. GE, Honeywell) but others have kept out (e.g. Phillips).

Consumer electronics: Sony, Panasonic, LG and Samsung have incorporated connectivity into their products. The TV is often used as a control hub. This strategy works when devices are from a single vendor as interoperability between different vendors is challenging.

Product specialists: This segment includes manufacturers of different home products such as locks, alarms, sensors, garage door controllers, and other products. August, Big Ass Fans, Kidde, Rachio, Schlage, Skybell, Yale are examples of this segment. Product specialists focus on incorporating wireless connectivity into their products and integrating with a multiple home gateway vendors.

Service providers: This segment includes connectivity service providers (mobile and fixed access service providers) as well as monitoring specialists like ADT and Vivint who offers their own home automation systems. Connectivity service providers developed home automation products in partnerships with product specialists. The business model is based on recurring fees, for example, AT&T Digital Home allows monitoring of security and energy starting at \$5/month. The trend is for service providers to become a one-stop-shop for home automation devices, gateways and Cloud service as exemplified by KT, NTT DoCoMo and PCCW.

¹HP, "HP Study Finds Alarming Vulnerabilities with Internet of Things (IoT) Home Security Systems," February, 2015.

²Symantec, "Security Response: Insecurity in the Internet of Things", Mario Ballano Barcena and Candid Wueest, Version 1.0, March 12, 2015.

Startups & peripheral vendors: This segment comprises home automation solution vendors who have sprung up especially within the last 3 years. This group focuses on developing a home hub and a few complementary devices most sought out by customers. They leverage partnerships with product specialists to provide a broader range of connected devices. Interoperability is critical to this approach. Startups typically seek to support many technologies in their home hubs to broaden their appeal.

Retailers: Examples of this segment include Lowe's Iris system to control security cameras, light switches, locks and other devices; and office superstore Staples offers a similar system called Connect. These vendors have to compete with the well-known technology brands and their long-term presence in the market will be tested.

Semiconductor vendors: ARM, Intel, Qualcomm and others are active participants in the home automation ecosystem which is a vehicle to drive semiconductor sales. Qualcomm is leading activities at the AllSeen Alliance for interoperability of devices. These companies can make investments into product companies as exemplified by Intel's acquisition of wearable health-tracking device company Basis Science for \$100 million in March 2014.

Many types of players constituting the home automation ecosystem leads to a complex channel to reach the end user. System integrators, device manufacturers, connectivity or Internet service providers, home automation system vendors, can reach the end user through a number of channels such as retail, direct, or through a partnership with another member of the ecosystem.




Figure 1: Home automation market value chain.

Emerging Trends

The adoption of wireless technologies for connectivity has shaken up the home automation market unleashing a number of trends:

The race to own the gateway. Home automation devices connect to the Internet through a gateway in the home. A standalone hub, cable set-top boxes, xDSL routers, a tablet may all serve as potential gateways. Established and startup companies are in fierce competition to own the gateway: Insteon (Microsoft), Nexia, Revolv (Google), SmartThings (Samsung), VeraLite, Wink (GE) are a few examples. The gateway owner furthers the chances of its technology platform, increases hold on the user and potentially gains access to a wealth of information on user behavior to derive additional revenues.

Proliferation of technologies. There are many wireless technologies used in home automation including Bluetooth, Wi-Fi, ZigBee, Z-Wave, and proprietary protocols. Technology is related to weighing tradeoffs against application requirements which define parameters such as range, power consumption, reliability, data rate, security, and addressing. Home automation companies are forming alliances and forums to ensure interoperability, improve user experience, and expand market power. This has led to clustering of large industry players jockeying for supremacy in different camps.

Building intelligence. Data science and machine learning techniques could be applied to user data to derive information, such as predicative behavior, leading to differentiated services and new revenues. This multifaceted issue remains nebulous at this stage as it has implications on consumer privacy that many regulators are grappling with.

Heightened competition. Do-It-Yourself (DIY) kits and luxury installations are creating competition across previously separated home automation segments. In parallel Cloud-based services and general-purpose controllers are driving market growth.

Scramble for security. Connectivity of home automation systems to the Cloud expanded the risk of attack and the compromise of privacy. Consequently, solution providers are jockeying to address security and privacy flaws.

Reliance on the Cloud. The Cloud can be used for storage, compute and networking which serves to reduce the cost of user devices and ease the introduction of new services. Leveraging the Cloud requires a framework for management, storage and backup, development of SaaS model, and interface with private and public service providers. Integration of Cloud services is changing the way peripheral devices are built as functions are moved to the Cloud leaving questions on interoperability.

Alliances & partnerships. As industry players converge on this market from different vantage points, they are forming alliances and partnerships to better capitalize on the opportunity, not to mention a heightened phase of M&A activities.

Emergence of application layer standards. The fragmented and siloed nature of home automation applications led to the emergence of application layer standards to allow multiple devices based on different technologies to interoperate and share data in a manner useful to the end user. This is exhibited by efforts undertaken by organizations such as the AllSeen Alliance and Open Interconnect Consortium.

New business models. The extension of capabilities brought about by connectivity and interworking of devices and between devices and the Cloud, is opening new avenues to price services and leading to new alliances. An example is the partnership between thermostat vendors and energy companies to offer rebates on new thermostat for users who exchange power consumption data with the power company.

Additionally, opportunities are brewing for residential IoT partnerships. For example, Pebble, a wearable company focusing on watches, has apps to control the Philips Hue lighting system and the Wink home automation system³. Even automotive companies such as Daimler are partnering with Nest Labs⁴ – their proof of concept to connect a Mercedes-Benz to a Nest thermostat provides an M2M connection of two consumer products.

The rapid changes in the home automation landscape is forcing companies to change their strategies to adapt to new realities through overhaul of product lines, establishing new partnerships, investing in new markets or making acquisitions, and developing new go-to-market strategies. After Belkin acquired Zensi in 2010⁵, Belkin launch their WeMo home automation system two years later at CES, to increase its relevance to consumers, taking the silo/gateway approach.

Market Evolution

Connected-home device shipments are projected to grow at a compound annual rate of 67% over the next five years, much faster than smartphone or tablet device growth, and hit 1.8 billion units shipped in 2019⁶. HVAC and security segments, including devices like connected thermostats and smoke detectors, will become popular first, leading the way to broader consumer adoption. This category makes will make up about 27% of shipments within the broader Internet of Things market in 2019 from about 25% today.

Market revenue is expected to reach over \$22.5 billion in 2018 from about \$12 billion in 2013 (CAGR of 13.7%). North America (31%), Europe (29%), advanced Asian economies (Japan and Korea), in addition to India and China (15%) are expected to lead the market.

³*How to Control Your Smarthome with Your Pebble Smartwatch*

⁴ *Mercedes-Benz at the 2014 Consumer Electronics Show: The Future Starts Now*

⁵ *MIT Technology Review. 'Home Sensor Startup Snapped Up. April, 2010.*

⁶ *November 2014. According to Business Intelligence, connected-home devices include all smart appliances (washers, dryers, refrigerators, etc.), safety and security systems (internet-connected sensors, monitors, cameras, and alarm systems), and energy equipment like smart thermostats and smart lighting.*

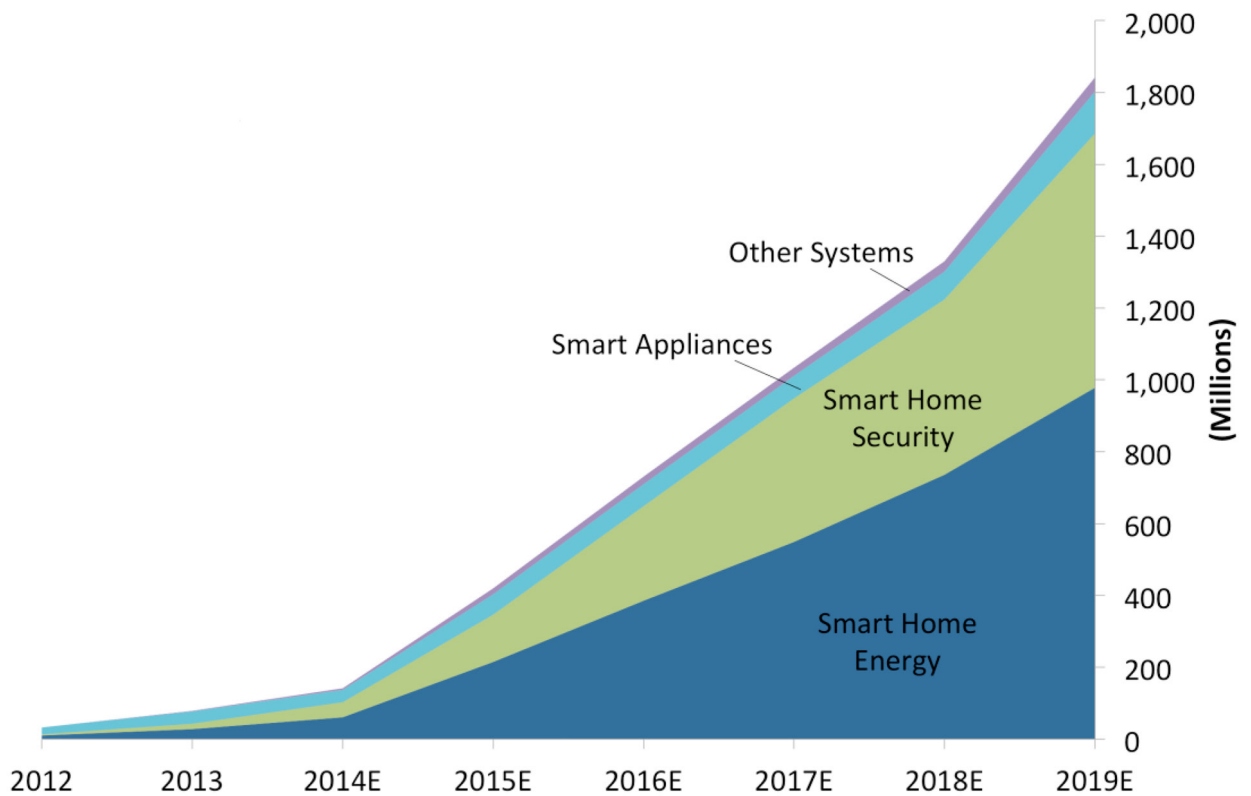


Figure 2: Global connected-home device shipments. [Source: Business Intelligence]

Consumer awareness and interest in connected-home devices is growing significantly. In the US, nearly two-thirds of broadband-equipped households are interested in a connected-home device bundle from their wireless service providers, according to survey data from Parks Associates. Millennials and people who have been in their home for between 3 – 4 years are the most inclined to buy connected-home devices. In each of these demographic groups, 10% of US residents already own a smart home device⁷. American consumers have ranked security as the highest benefit of home automation ahead of convenience and savings⁸.




Figure 3: Security is the top benefit for half of Americans. [Source: MaRS Market Insights]

⁷ Business Intelligence. "The Connected-Home Report: Forecasts and growth trends for one of the top 'Internet of Things' markets," March 2015.

⁸ MaRS Market Insights, "The Connected Home: Smart automation enables home energy management," October, 2014.

Market growth projections are matched by an increasing rate of investments. Smart home startups took \$454 million in investor funding in 2014, an increase of 57% over 2013⁹. Among the largest deals in the space over the past six months include a \$38 million Series B from Bessemer Venture Partners, Comcast Ventures and Qualcomm Ventures to August (smart locks) and a \$31.8 million Series B to connected home software platform Zonoff from investors including Grotech Ventures and Valhalla Partners.




Figure 4: Home automation funding trends.
[Source: CB Insights, Xona Partners Estimates]

The home automation market is an active field for M&As as the boundaries between different ecosystems are blurred in the drive to capture market share with large companies placing early bets through acquisitions. Moreover, some of the lesser known companies have matured their businesses with Control4 and Alarm.com completing successful IPOs in the past two years, in 2013 and 2015, respectively. While both companies took over 10 years to achieve this, these durables goods companies are showing that there is still a massive opportunity to convert consumers.




Figure 5: Home Automation Acquisitions 2013 – 2015.

⁹ CB Insights, "Disrupting Honeywell: The Startups Unbundling Honeywell in the Smart Home," April 2015.

Conclusions

The home automation market is in the midst of a rapid change which began a few years ago at the advent of smartphones and wireless data service. It has since accelerated to breakdown the boundaries of established players as new entrants challenge traditional approaches with new business models, technologies and applications. The ecosystem continues to expand with record high investments and M&A activities. The race to own the home automation market is pursued from different angles by companies in adjacent sectors. The result is an ecosystem at an early stage of formation as alliances are beginning to coalesce to meet the needs of ever more knowledgeable and demanding users. The home automation market will remain a focal point within the greater IoT market space as applications expand as does the pressure to consolidate which will usher a new phase of market development.

XONA PARTNERS

and

OLSWANG

Internet of Things

Roadmaps and Regulatory Considerations

Dr. Riad Hartani, Frank Rayal, Rolf Lumpe
(Xona Partners)

&

Purvi Parekh (Olswang)

October 2015

Preamble

The Internet of Things (IoT) is by definition a vast topic that encompasses multiple markets, technologies, and disciplines. IoT comes with the promise of a new wave of applications and services deployment, significant investments and returns. Together with such promise, comes a series of obstacles – commercial, technical, regulatory and legal - that combined could slow down the rate of adoption of many smart technologies. IoT applications are broad, fragmented and (currently at least) siloed in specific verticals where multiple competing technologies (and law) vie for prominence. The topics of security and privacy become complex. Questions around the adequacy of resources for M2M services are paramount. Consumer acceptance of M2M services is fundamental.

From this perspective, IoT is an evolutionary process that will exhibit varying adoption rates in each silo while the market and regulators work their way through the challenges.

In this paper, we set out an ecosystem reference model for IoT and provide a brief overview of some key challenges, with special emphasis on the legal and regulatory aspects, how they are being addressed and how upcoming changes may impact in the future.

The IoT Ecosystem

To conceptually define IoT, consider a five-layer functional model that includes devices, connectivity, applications, platforms, and services (Figure 1):

Devices: Sensors, identifiers and gateways are types of IoT devices used to collect and convey information. Devices are designed and deployed to meet the application use case requirements. They can range from simple identifiers that provide specific information on the object, to complex devices that have the ability to measure (sensors) and process data (gateways). The application, use case and deployment scenario places requirements on the device such as size, weight, power consumption, and life of operation or deployment. This in turn impacts the connectivity of the device to the network. A variety of IoT devices have emerged in various business verticals, starting in the utility / energy sectors and evolving to devices in the health, transportation, home and finance ecosystems amongst others.

Connectivity: Devices can be connected directly to the network, or indirectly through another similar device (mesh) or a gateway that is provisioned to support multiple devices. Connectivity can be through a number of physical media such as copper, fiber and optical cable, or through the air through a number of wireless technologies. One of the challenges in IoT is the proliferation of connectivity standards, which is a common symptom of the breadth and fragmentation of IoT application requirements. These standards span the entire logical protocol stack through layers 1 – 7. Examples of connectivity would include the traditional 2.5/3/4G networks, as well as various local area solutions (Zigbee, Wi-Fi, Bluetooth, others) and low power wide area solutions (e.g. Weightless) among others.

Applications: Applications define the use case of the device and include all the necessary functions required to make use of the device for the intended purpose including the hardware and software architectures. IoT application stores are emerging with applicability to specific industry verticals, with the health wearable devices being a recent example.

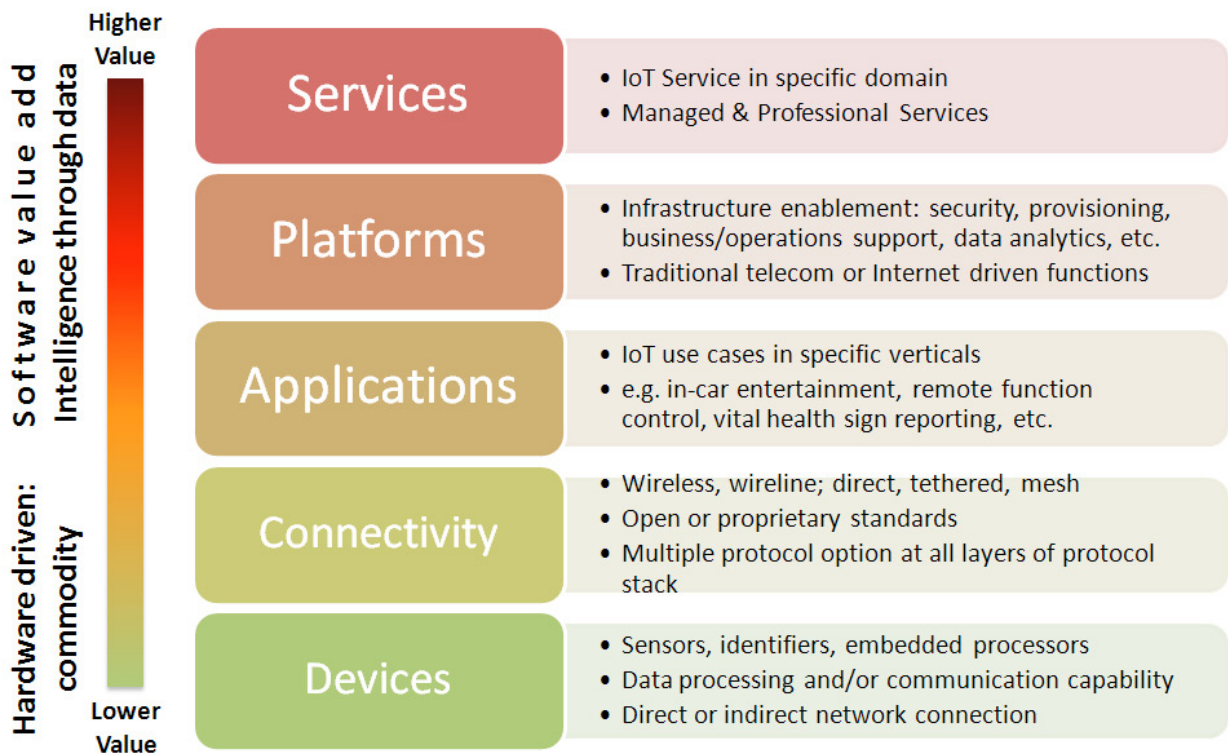


Figure 1: IoT ecosystem reference model.

Platforms: Devices and connectivity require a platform to provide a service. Platforms are used to provision devices, manage and control them. They are used for billing and fraud detection. Platforms also provide the means to customize functions and data according to the requirements of end users. From this perspective, platforms allow the IoT infrastructure to perform as required.

Services: This references the IoT service to the end-customer. The service provider leverages all the downstream elements in this value chain: platforms, applications, connectivity and devices. The service provider can be the same or different from the platform and application provider. Examples include automotive automated diagnostic, medical geriatrics and remote power consumption optimization.

The IoT Connectivity Model – The data

In order to put it into context our conclusions and observations on IoT development, we model data flow, which can be characterized by three stages: data creation, transmission, and consumption.

Data creation: Data is generated by different types of devices, Data has specific characteristics such as rate, volume, latency, and frequency. For example, video surveillance has a high data rate whereas Supervisory Control And Data Acquisition (SCADA) systems have a low bit rate. Taking this example further, in many SCADA applications, the latency has to be very low to accommodate specific requirements of an application such as a fault in an electric transformer that requires the instantaneous switching of electric currents to avoid damage while there is higher tolerance to latency in video applications.

The creation of data can bring with it data privacy and security concerns at both a user level and a regulatory level. Although data flows may appear small, they still leave a digital trail. By the

same token whilst a specific silo of information may appear harmless, putting silos together can provide a detailed insight into a person's life, opening them up to user profiling or tracking. An added complication is the different layers in the privacy evaluation; data is not just recorded in the database of an M2M service provider, but also in the database of the mobile network provider and/or in a home gateway or device. Add cloud services into the mix and the locations and jurisdictions where data resides also increases. All of these factors bring IoT and M2M services into the realms of data privacy legislation.

From a policy perspective the regulatory approach on IoT has not favored the creation or adoption of bespoke IoT legislation. The reality is that there is plenty of vertical legislation that applies to the IoT ecosystem under communications, privacy and sector specific laws, much in the same way as it applies to the majority of new technologies in the market. Instead a favored approach is that of "privacy by design" i.e. taking privacy and human values into account throughout the whole IoT engineering process. The concept, which actually originated a decade¹ ago, has been given the regulatory thumbs up across both sides of the Atlantic while Asia is closely observing the adaptation. Recently Federal Trade Commission Chairwoman Edith Ramirez endorsed the idea of companies conducting privacy (and security risk) assessments during the design process as well as the testing of security measures before products launch. Her endorsement went wider than just the engineering phase; she was also supportive of ongoing monitoring of products for vulnerabilities throughout their life cycle.

Data transmission: The transmission of data raises questions around bandwidth, latency, compression, encoding, multiplexing and security, especially when considering the various platforms and networks over which data may traverse. Data encryption and device authentication are commonly adopted to combat security concerns. In addition, and although not mandated by regulators, commercial contracts in the IoT value chain increasingly incorporate detailed provisions around security defining responsibilities and liabilities as between all of the parties in the IoT value chain – not just the two parties at sitting at the negotiating table. These provisions range from stringent obligations on protecting against false requests for information to implementing ways to identify and combat unauthenticated commands. User behavior is also legislated for, with users being mandated to change passwords at regular intervals.

As with data creation privacy remains a concern, particularly where data is being transmitted across different countries or being routed to countries which do not have the same level of data privacy protection as exhibited in the country of origin. Data protection rules already tailor for such transfers and how these are to be handled in order to safeguard its protection.

Data consumption: Data is consumed in different ways, depending on the application. Simple systems that involve the user directly interacting with device is a mainstream medium. Think of the interaction with a wearable through an application on a mobile device or tablet. Alternatively and increasingly, sophisticated techniques based on data sciences are used to seek information beyond the original intended use. Whilst the collectors of that data promote the benefits that such data collection could result in (e.g. a homeowner may install a Google Nest thermostat, which she can control remotely; however, the data can also be shared with the utility company to control temperature within certain bounds during peak hours and to create more overall efficiencies),

1. Joint report on "Privacy-enhancing technologies" by Information and Privacy Commissioner of Ontario, Canada, the Dutch Data Protection Authority and the Netherlands Organization for Applied Scientific Research

from a privacy perspective there are immediate concerns. Users do not want to be tracked or profiled unless they have specifically consented.

Observations on IoT Market Space

By deploying the conceptual IoT framework above, we can model developments across the ecosystem layers starting with devices and connectivity and ending with platforms and services.

To start, we note that the IoT use case requires devices and connectivity, underpinned by the interoperability of services, devices and platforms. Device characteristics such as size, weight, placement, mobility, power and communication characteristics as defined by the application drive what connectivity is required. Each vertical market (for example, automotive, utility, agriculture, home, health, general industry, etc.) uses different options thus resulting in a proliferation of connectivity standards. Whilst there are attempts at harmonization and standardization across verticals, we are not yet in a place where it is the norm.

1. Proliferation of connectivity standards: Depending on the characteristics of connectivity, various standards have been, or are, in the process of being defined. 3GPP standards such as GPRS, UMTS and LTE are licensed band access schemes that rely on high power for long range, consequently are relatively expensive in comparison with other connectivity techniques. On the other hand, technologies such as Bluetooth are meant for short-range communications in unlicensed spectrum and are low on power consumption. Various LPWA proprietary solutions have also recently emerged, mostly in unlicensed sub-1GHz spectrum but also in some licensed bands. Wi-Fi relies on higher power and provides longer range than Bluetooth albeit at a higher cost.

In recent years, advancements in silicon technologies such as 28 and 14 nm processes have significantly reduced power consumption to allow ever-smaller devices with less battery requirements to come to market. Coupled with the maturity of smartphones, this has led to the significant increase in wearables and personal connected devices.

From a regulatory standpoint international adoption through common standards has been on the agenda of many regulators and interested stakeholder bodies, keen not to stall the advancement of IoT. Only a few weeks ago, IEEE, the world's largest professional organization dedicated to advancing technology for humanity, announced that the Industrial Internet Consortium® (IIC) and the IEEE Standards Association (IEEE-SA) were collaborating toward development of a comprehensive architecture for an interoperable Internet of Things (IoT) around the world. In parallel various verticals are looking specifically at better harmonization. Take the automotive sector for example where new legislation was announced in the US around the creation of federal standards that secure cars and protect drivers' privacy.

2. Commoditization of devices: Essential to enable the business case for IoT applications is the trend of cost reductions in devices, as illustrated with the large number of players commercializing consumer wearables (Figure 2). The challenge to device manufacturers is how to differentiate from competitors. Our observation is that software applications and platforms, including operating systems, are the essential leverages used by device manufacturers to differentiate (e.g. Apple/iOS, Google/Android; Samsung attempt at differentiating through Tizen, and in a similar way Alibaba and Xiami's own platforms design).



	Xiaomi Mi Band	Fitbit Flex	Jawbone Up
Steps taken	✓	✓	✓
Calories burned	✓	✓	✓
Distance traveled	✓	✓	✓
Active time	✓	✓	✓
Sleep time	✓	✓	✓
Sleep quality	✓	✓	✓
Map routes		✓	
Average pace		✓	
3rd party apps		✓	✓
Diet tracker			✓
Waterproof	✓	✓	✓
Bluetooth	✓	✓	
Social sharing	✓	✓	✓
Alarm	✓	✓	✓
Notifications			✓
Indicator lights	✓	✓	✓
Price	\$13	\$99	\$79

Figure 2: Device commoditization.

3. Commoditization of connectivity: As with devices we are seeing a downward slope of cost reduction for connectivity costs of IoT applications, driven by the need to enable the business case of most applications. There are many variants of connectivity including wireline and wireless technologies and increasingly a spectrum in between of license free/license exempt wireless opportunities. The lowest cost wireless connectivity leverages license-exempt spectrum over short distance (Figure 3). Wearables, for example, leverage Bluetooth to connect with smartphones. Alternatively, some consumer devices rely on longer-range license-exempt technologies such as Wi-Fi. Central hubs for connectivity and routing are deployed to tether over longer distances for remote control and monitoring. Where mobility is required, wireless technologies in licensed spectrum can be implemented albeit at a higher cost. It is exactly because of this that regulators are looking to support the IoT business case by considering comparable spectrum solutions that fall within the spaces between the licensed bands.

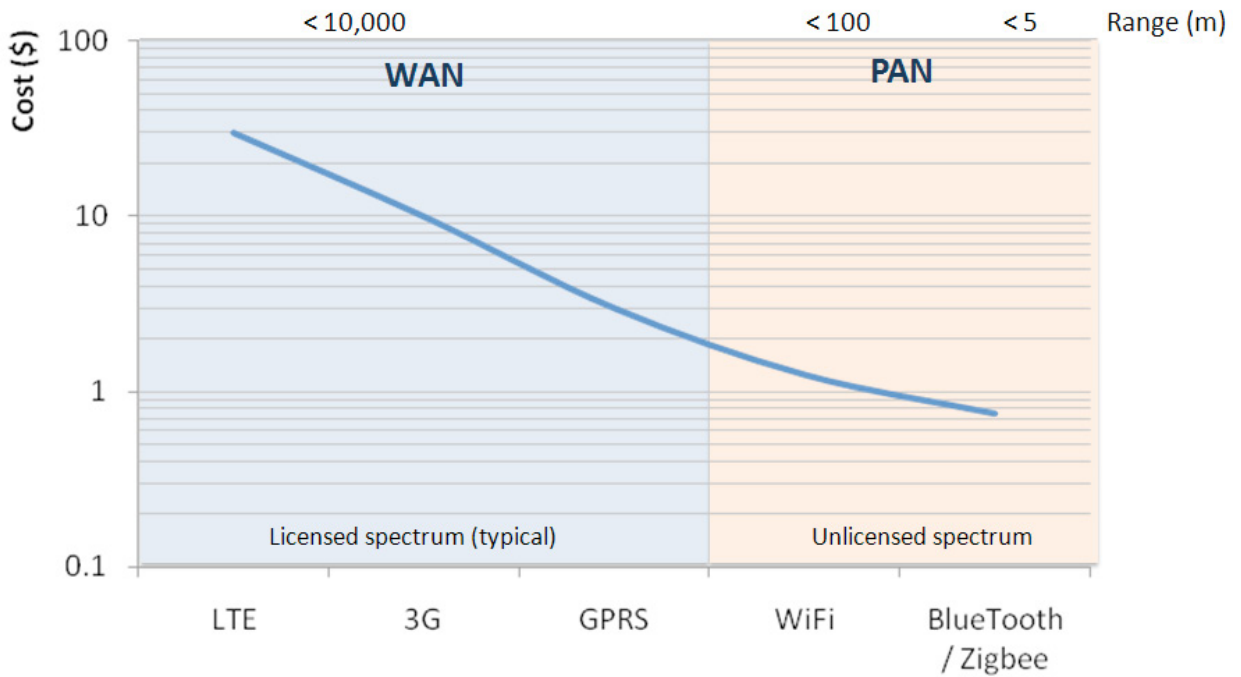


Figure 3: Cost dynamics for IoT wireless connectivity.

4. Emergence of long-range low power wireless technologies: We see an opportunity for very long range wireless technologies that are low power, low cost and work over long ranges (Figure 4). Such technologies are now on the market but it is still early days in the proof of their commercial viability (for example, Neul Weightless, SigFox Ultra Narrow Band, Semtech LoRa, and On-Ramp). These technologies often assume the build-out of a parallel IoT network to the mobile network.

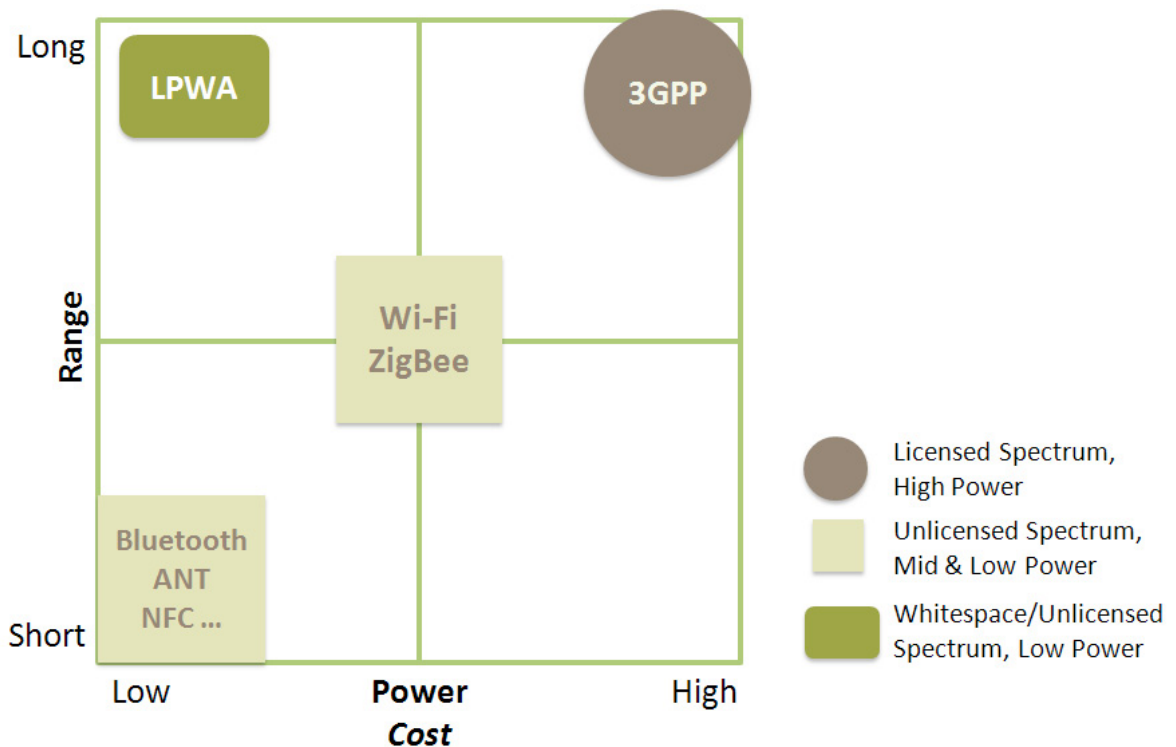


Figure 4: IoT Wireless connectivity.

5. Competition and harmonization of connectivity standards: Connectivity standards have been progressing slowly but steadily. The challenge is not in the definition of these standards, but more in the number of variety of competing and complementary standards, as well as the conflicting interests of the industrial groups behind the various standards. Although harmonization is ongoing, it is very likely that IoT solutions will face challenges for rapid mass adoption. The development of interworking platforms with open APIs will help alleviate some of these challenges by allowing interoperability of different standards or different implementations of the same standards. This is not only the case for physical and link layer standards, but also includes aspects related to applications and services running on top of the IoT ecosystem.

6. Partnerships and alliances to win the IoT platform war: The development of IoT solutions is inherently about the development of ecosystems around offered solutions. Such ecosystems are not mandated by legislation but instead built via negotiated partnerships between various industry players. Given regulatory challenges on revenue, the leading players will seek to control the ecosystem by providing a platform that would host IoT applications, and over which IoT services will be built (Figure 5), as this is an important new revenue stream for them. As in any platform model, such as those in smartphones and the Internet, the key is to increase its adoption. Various models are being put in place to achieve this, via the development of open source IoT connectivity and interworking software, open APIs to plug into the platforms, and SDKs to develop services on top of the platform. We foresee the emergence of selective alliances over the next few years, across industry verticals, with a focus on advancing specific IoT platforms, but progressively evolving towards selecting winners, as it's traditionally the case for Internet-centric business models. Various contenders are already in the game to achieve this, including the Internet platform players (Google, Apple, Amazon, etc.), the lead industrial players with a specific vertical focus (e.g. GE for industrial Internet), as well as mobile operators, particularly those who support a strategy towards Internet-scale OTT deployment.

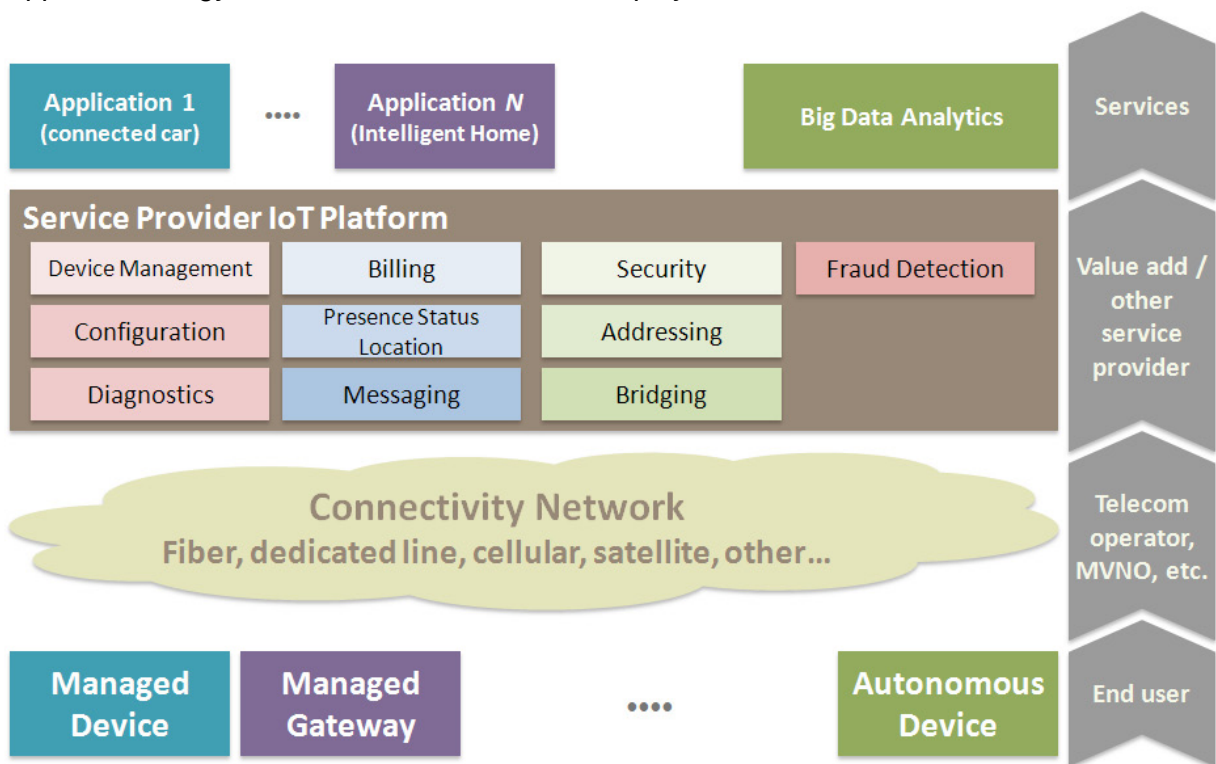


Figure 5: Value appropriation through platforms.

7. Emergence of new MNO and MVNO service models: A key dynamic of the IoT market is that the majority of ‘value’ in any IoT application lies not in the simple carriage of data, but in the provision of an overall service. For example, a wide-area wireless enabled home security system represents a significant revenue opportunity for a mobile or virtual mobile network operator, including revenues from device sales, installation, and monthly service fees. However, the data traffic revenue that such a solution generates is likely to be relatively small in comparison. Similarly a connected health solution will include the connectivity network as well as the platform to manage the solution, interfacing with the various stakeholders in the health solution value chain. The story is the same for many other IoT applications: the real opportunity for mobile operators lies in moving up the value stack and away from the simple provision of data carriage services. The result is an ecosystem that is complex, multi-party and heterogeneous in nature. The mobile network operator provides the connectivity and IoT management for value add. IoT solution providers (either OTT service providers or mobile service providers who offers IoT solutions) will have to integrate all the components of the ecosystem for the end-to-end IoT solution. Each has a crucial role to play in the value chain.

With such a complex chain comes a mesh of legal liabilities and challenges. Individual responsibilities are more difficult to segment and hence to legislate for. Privacy and privacy responsibilities require more careful scrutiny as do security requirements and obligations around where these sit in the value chain. Commercial leverage will play an even more important part.

8. Extracting value through data sciences: Temporarily putting to one side the fundamental privacy issues around extracting value in this way, as businesses evolve to leverage the huge amounts of data assembled, mining and learning through such data creates significant opportunities. By the same token so does optimizing communication between those producing it and those using it. The desired goal of IoT businesses is to create a solid foundation architecture that is able to provide these optimal functional capabilities together with a platform to overlay data science applications. This would include the various layers in the data value chain – optimized processing through an acceleration of migrations to the cloud, scalable data management leveraging big data models and the use of customized data sciences solutions for business intelligence creation. This “solution” is complemented by a fundamental re-architecture of IT models within the businesses integrating IoT models.

We are now witnessing the emergence of an enhanced (and new in some cases) set of machine learning and data mining algorithms, specifically focused on clustering and predictive modeling in high dimensional spaces which is based on imprecise, uncertain and incomplete information, efficient statistical data summarization and features extraction algorithms as well as large-scale real-time data stream management. These tools will be at the core of the processing engines being commercialized or running in open source environment, and will aim, when applied to specific industry problems, at optimizing the existing business logic and augment it with new functionalities over time.

9. Evolution to 5G: In terms of timing and mass market adoption of advanced IoT solutions, it is very likely that this will converge and overlap with the specification and rollout of the first 5G networks. It is then natural that 5G specifications will have to take into account IoT requirements, either directly or via the complementary technologies that will form the future mobile ecosystem (including evolutions of Wi-Fi, LPWA, Zigbee, etc.). As such, the LTE roadmap will continue to evolve to include new features that represent a precursor to those in 5G. For example, LTE-

MTC in Release 13 aims to reduce power consumption of LTE devices for IoT applications and achieve low cost points by eliminating some of the broadband features of LTE (Figure 6). On the core, back-end and underlying IT infrastructure, a gradual move towards virtualization, specific functionality enablement in private/hybrid/public cloud environment, and integration of big data analysis frameworks into network data management, will start appearing. All of these aspects will in essence contribute to bringing advanced IoT solutions and IoT centric business models to markets.




Figure 6: IoT ecosystem dynamics.

IoT – The Road Ahead

As far as mass adoption is concerned the IoT era has had various false starts. The recent convergence of various trends including innovation in low power and low cost device technologies, scalable network connectivity as well as mainstream cloud and big data processing models have opened a new window for the emergence of IoT based value added services that will in time become mainstream. Vertical specific creation of common standards, legislations and regulatory approach will further support greater international deployment.

With the significant transformation in the IoT ecosystem comes challenge and opportunity. IoT, in its blurring of the distinction between public and private, is driving increased focus change in the business and legal landscape, with significant implications for regulatory and policy makers over the next decade. Indeed, as most recently highlighted by BEREC's (the Body of European Regulators for Electronic Communications) report and public consultation on M2M services, this focus is about how to facilitate M2M and to make it thrive. Whether this means that we are looking at more rather than less regulation remains to be seen.

Acronyms

3G	Third generation
3GPP	Third generation partnership project
4G	Fourth generation
5G	Fifth generation
API	Application program interface
DSL	Digital subscriber line
GPRS	General packet radio service
GPS	Global positioning system
GSM	Global System for Mobile communications
iOS	iPhone operating system
IoT	Internet of Things
IP	Internet protocol
IT	Information technology
LPWA	Low power wide area
LTE	Long Term Evolution
MNO	Mobile network operator
MTC	Machine type-communication
MVNO	Mobile virtual network operator
P2P	Peer to peer
PLC	Power line communications
SCADA	Supervisory control and data acquisition
SDK	Software development kit
SIM	Subscriber identity module
UMTS	Universal Mobile Telecommunications System
V2P	Vehicle to Pedestrian communications
V2V	Vehicle to Vehicle communications
WRC	World Radio Conference
IP	Internet Protocol
IS	Industry Standard
ITU	International Telecommunication Union
JDBC	Java Database Connectivity
LTU	Long Term Evolution
M2M	Machine to machine
METIA	Mobile and wireless communications Enablers for Twenty-twenty (2020) Information Society
MIMO	Multiple Input Multiple Output
MME	Mobility Management Entity
MTAS	Multimedia Telephony Messaging Server
MVNO	Mobile Virtual Network Operator

XONNA PARTNERS

Shaping Cellular IoT Connectivity

Emerging Technologies in
Wide-Area Connectivity

Frank Rayal

June, 2015

Table of Contents

1	INTRODUCTION	3
2	RECENT DEVELOPMENTS IN CELLULAR DEVICE CONNECTIVITY	3
3	LTE IoT CONNECTIVITY	5
4	LPWA IN LICENSED BANDS	8
	4.1 Semtech – LoRa Technology (Long Range)	8
	4.2 SigFox Cooperative Ultra Narrowband (C-UNB)	8
	4.3 Huawei/Neul	9
	4.4 Qualcomm – NB-OFDMA	9
5	THE IMPLICATIONS	10
6	CONCLUDING REMARKS	14
7	ACRONYMS	15
8	THE XONA PARTNERS TEAM	16

1 Introduction

Creating networks of things is widely considered as the next engine for economic growth valued in the trillions of dollars. Yet creating the Internet of Things (IoT) is not a trivial activity as demonstrated by inflated expectations that have been slow to materialize as anticipated by market research analysts. The IoT market remains highly fragmented with multitudes of applications, each with its own set of requirements that adds constraints on the type of connectivity solution. While connectivity is only one element of the IoT ecosystem stack, it is a prerequisite to all other layers for without connectivity, IoT would not exist. From this perspective, IoT can only take off with the availability of cost effective connectivity solutions that meet both business case and the technology requirements of the applications.

One segment of IoT revolves around wide area connectivity of devices. Cellular technologies such as GPRS and 3G UMTS dominate this market today. Where these technologies have proved expensive, mesh solutions are used to create wide area networks based on relatively short connectivity segments. Satellite is used in remote areas where the business case works. In this paper, we discuss the emerging low-power wide-area (LPWA) connectivity technologies which have unique characteristics as they are purposely designed to meet wide-area IoT application requirements unlike the other technologies which are adapted for IoT. LPWA technologies are typically narrowband (with some exceptions) and operate in the ISM license-exempt spectrum bands. In recent months, GERAN and 3GPP standards organizations embarked on a process of standardizing narrowband technology for use in mobile spectrum. Several proponents of LPWA technologies have put forward their technologies. The competition in the standards race extends to 3GPP, where the roadmap for cost reduced LTE module for IoT applications is under development (LTE-M), and other standard organizations that are focusing on 5G technologies.

This paper is divided into two parts. The first is focused on technology where we provide an overview of narrowband LPWA technologies. We also discuss the roadmap for LTE-M to compare and contrast the solutions. The review of technology allows us to better understand the implications strategy, markets and ultimately the potential success of each approach. In the second part of the paper, we present a discussion on evolving market dynamics where high stakes are in play to determine the winners of the next round of market growth drivers.

In the context of this paper, we define ‘device’ as a connected object that excludes consumer electronics including smartphones, tablets, dongles, e-readers and such devices. We also use the term IoT instead of machine-to-machine (M2M) connectivity which is traditional in industry circles because we seek to emphasize an encompassing value proposition beyond connectivity.

2 Recent Developments in Cellular Device Connectivity

Cellular device connectivity constitutes a relatively small fraction of total connected devices – estimated at 243 million in 2014, or about 3.5% of total connected devices. The vast majority of these devices, 77%, use 2G GPRS which is a technology first commercialized in 2000¹. The cost of 2G modules have dropped in recent years to reach about \$10/module in volume while the cost of LTE modules are around \$50. By 2020, 1 billion cellular connected devices are expected with 2G accounting for 44% of connectivity while 3G and LTE will account for 33% and 23%, respectively.

¹GSMA Intelligence, “Global cellular M2M technology forecasts and assumptions,” March 2015.

Applications of cellular connectivity remain concentrated in traditional applications such as transportation, automotives, and location management. Cellular 2G connectivity provides the benefit of world-wide coverage and almost-unified frequency spectrum allocation (900 MHz most of the world, 850 MHz in North America). The Embedded SIM technology (eUICC) simplifies the process of providing service through different operators which enables the mobility market. Nevertheless, there are limitations to cellular connectivity which LPWA addresses. These limitations fundamentally center on two key issues: high power consumption that does not allow battery operation over an extended period of time reaching into the years, and the cost of service which includes the cost of the device and the supporting infrastructure that factors into the return on investment for the service provider. The result is a bifurcation of wide area IoT technologies along three axes:

2.1 *LTE evolution:* LTE is fundamentally a technology for broadband connectivity. It was not designed to address connected devices. LTE consumes too much power and offers much higher capacity than required by many IoT applications. The modems are relatively expensive to integrate but into high-value applications with a good power supply such as a vehicle. The 3GPP standards body is addressing the shortcomings of LTE in IoT connectivity by incorporating enhancements in network access and defining new device categories that consume less power and reduce module cost by eliminating many of the broadband features such as multiple transceivers and antenna systems. New device categories include Category 0 (Cat0) which is defined in 3GPP Release 12 and sub-Cat0 which is in the process of being defined.

2.2 *LPWA technologies – unlicensed band:* Designed to cater to wide-area IoT connectivity, these technologies feature a protocol stack optimized for device access which typically consists of short messages sent in bursts. The physical layer is typically kept simple with low modulation scheme for robustness and low complexity. The medium access control layer is efficient with low overhead signaling in low data-rate, low network access periodicity use cases. LPWA technologies are designed for scalability on the order of thousands of devices per cell. They are deployed in license-exempt spectrum such as the ISM band (e.g. 902-928 MHz in North America, 866 – 870 MHz in Europe, 2400 – 2483.5 MHz world-wide). The LPWA market is dominated by startups and structured around verticals where two operational modes are emerging: private networks addressing a specific client, and public networks shared between different clients.

2.3 *LPWA technologies – licensed-band:* Although LPWA technologies are hardened against interference which is built into the protocol stack, licensed-spectrum operations enables greater assurance of reliability. Standardization coalesces focus on a technology, enables the creation of a wide ecosystem and improves economics. Availability of a standard gives service providers a greater incentive to enter the IoT market for new applications. For these reasons, standardization activities of narrowband LPWA technologies have begun at GERAN, the standard body responsible for GSM standardization, and has recently moved to 3GPP where 3G and LTE are standardized – a very significant development with high implications on wireless operators IoT roadmaps. Semtech, SigFox, Huawei/Neul, Qualcomm have put forward proposals to meet GERAN guidelines for narrowband IoT connectivity. We review these technologies later noting that there are some differences from the original unlicensed-band technologies in order to accommodate new requirements for compatibility with cellular networks operating in licensed spectrum.

The nascent LPWA market is set to disrupt the scene with mobile network operators taking different positions on how to address these upcoming technologies. Market forecasts for LPWA

vary between a low of 1 billion and a high of 3 billion connected devices by 2020, most of which in North America, Europe and the Asia Pacific region deployed in lead applications including smart cities, smart buildings, agriculture & environment, and utilities.

While LTE-M falls along the preferred roadmap for MNOs, its availability is later than LPWA technologies; and even when it becomes available, it would not meet all the requirements for wide-area IoT connectivity. LPWA in unlicensed bands represent a departure from the modus operandi of MNOs which revolves on licensed spectrum, reliability and personal broadband connectivity to which the core network and support systems are designed to for. Finally, LPWA in licensed spectrum appears as an attempt to harmonize the first two axes, but there are doubts that it would provide true differentiation from LTE-M, or even beat the timelines of LTE-M which may leave it with little market relevance. In fact, some contend that licensed-spectrum LPWA is a decoy against unlicensed spectrum LPWA. How the market will shape up in the coming months and years and what moves the different ecosystem players are making to assure a position in an emerging sector is beyond the scope of this paper. But we would provide some of the context for further analysis by covering essential elements of technology roadmap for LTE and a four LPWA technologies submitted for standardization at GERAN.

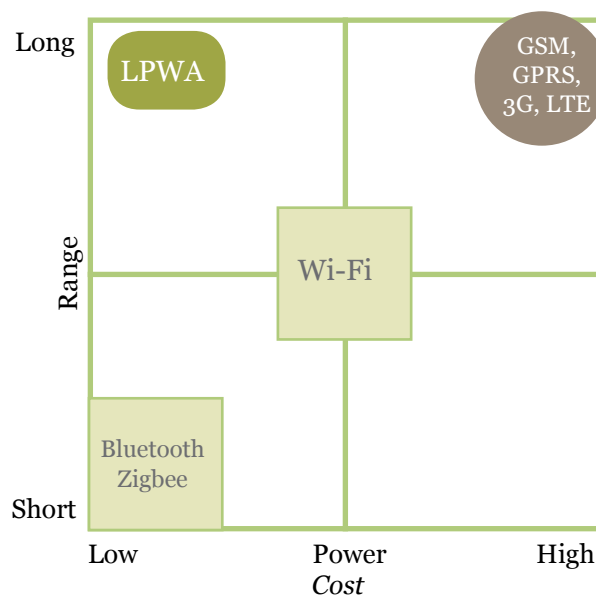


Figure 1: IoT connectivity technologies feature matrix.

3 LTE IoT Connectivity

The early LTE specifications defined in 3GPP Release 8 and 9 are focused on meeting requirements for mobile broadband connectivity in macro cellular network topology. 3GPP Release 10 first introduced the low access priority indicator (LAPI) to enable congestion and overload control mechanisms where the network can, for example, reject or delay connection request from low-priority devices in a congestion scenario. This is followed in Release 11 by incorporating architectural improvements that include the introduction of new functional entities for device connectivity (M2M-IWF and M2M-AAA) and eliminating the requirement for a phone number (MSISDN) in favor of IPv6 identifier.

LTE Release 8 through 11 presents several challenges for device connectivity:

- **Range:** insufficient system gain to reach deep into buildings and basements particularly for stationary devices.
- **Complexity:** multiple transmit and receive antenna configuration that is costly for IoT applications.
- **Scalability:** cannot support high number of devices which impacts the business case.
- **Power:** high power consumption does not allow operating on battery for extended periods
- **Inefficiency:** high signaling overhead in relationship to the amount of transmitted data for many applications.

LTE Release 12 begins to address these shortcomings by defining a new category of devices termed Category 0 targeted for device IoT connectivity. Some of Release 12 features include the following:

- One receive (Rx) antenna compared to a minimum of 2 Rx antennas for other device categories which reduces cost and complexity at the expense of losing diversity reception.
- Limited peak data rate to 1 Mbps in downlink and uplink in comparison with peak rate of 10 Mbps/5 Mbps in DL/UL for Cat1 device which is the lowest category of non-M2M LTE device. This is accomplished by reducing the transport block size.
- Optional half-duplex FDD mode that reduces the cost of the modem by eliminating a few hardware components (e.g. duplexer, switches).
- Enhanced Power Saving Mode (PSM). A device remains registered on the network but not reachable in PSM mode which eliminates registration setup and connection signaling. This optimizes modem turn-on for device-originated data or scheduled transmissions. It improves battery life and reduces overhead signaling.
- Extended Discontinuous Reception (DRX). DRX is designed for paging mobile user devices accounts for large amount of device power consumption. Increasing the DRX/paging cycle reduces energy consumptions by increasing the length of the sleep cycle but lowers device responsiveness which is acceptable in many IoT applications.
- Reduced Tracking Area Updates (TAU) and measurements for stationary devices.

While Rel-12 Cat0 device brings performance improvements for IoT applications, it is considered as a stepping stone for further improvements. Currently, a new device category is being defined as part of Release 13 specifications. It promises further reduction in complexity and cost by reducing the channel bandwidth to 1.4 MHz, lowering the data rate and reducing transmit power among other modifications to the protocol stack. It also targets improving the system gain by 20 dB over that for current 2G and 4G devices (typical maximum coupling loss of 140 dB) to over 160 dB maximum coupling loss.

Table 1 Feature list comparison for different UE categories. [Adapted from RP140845]

	Rel-8 Cat-4	Rel-8 Cat-1	Rel-12 Cato	Rel-13
Downlink peak rate	150 Mbps	10 Mbps	1 Mbps	~200 kbps
Uplink peak rate	50 Mbps	5 Mbps	1 Mbps	~200 kbps
Max number of DL spatial layers	2	1	1	1
Number of UE RF receiver chains	2	2	1	1
Modulation DL/UL	64 / 16 QAM	64 / 16 QAM	64 / 16 QAM	
Transport block size DL/UL (bits)	150752/51024	10296/5160	1000/1000	
Duplex mode	Full duplex	Full duplex	Half duplex (optional)	Half duplex (optional)
UE receive bandwidth	20 MHz	20 MHz	20 MHz	1.4 MHz
Maximum UE transmit power	23 dBm	23 dBm	23 dBm	~20 dBm
Modem complexity relative to Cat-1	125%	100%	50%	25%

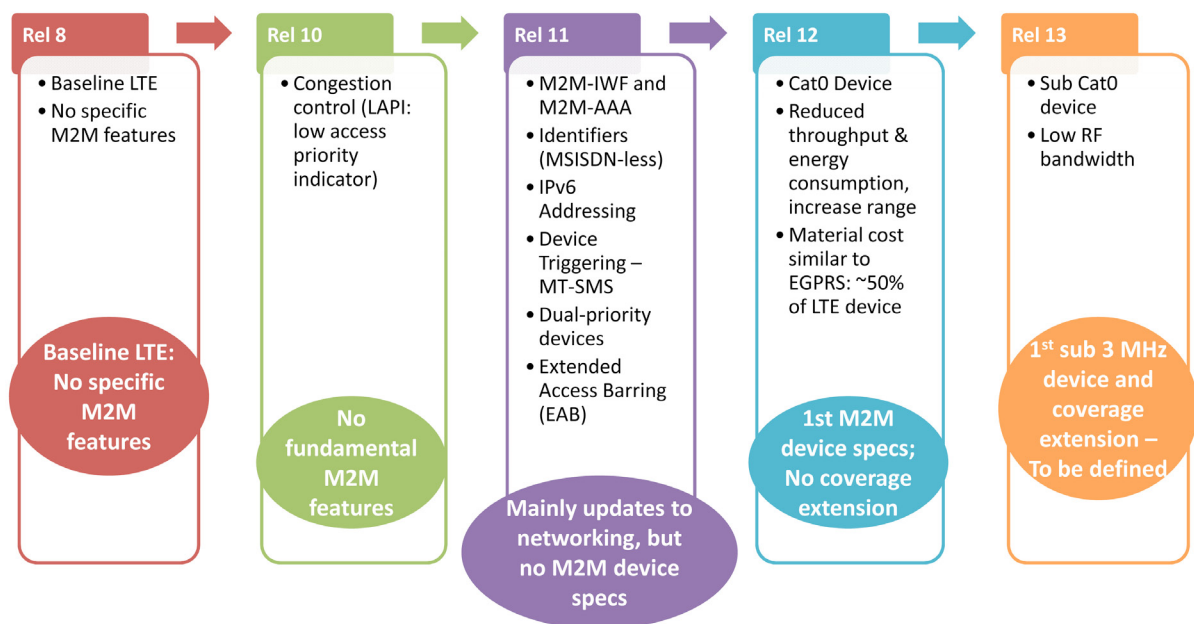


Figure 2: LTE roadmap to support machine-type communications.

4 LPWA in Licensed Bands

4.1 Semtech – LoRa Technology (Long Range)

The proposal by Semtech to GERAN revolves on adapting the current LoRa technology which operates in sub 1 GHz ISM bands. The LoRa technology defines two physical layer modes:

1. Narrowband mode targeted at fixed devices.
2. Chirp spread spectrum (CSS) targeted at mobile devices and devices embedded deep into buildings. This mode provides positioning information at the cost of lower spectral efficiency than narrowband mode.

Both physical layer modes operate in 200 kHz channel bandwidth similar to GSM. In the narrowband mode, the uplink is divided into 72 channels of different bandwidth ranging from 400 Hz channel placed at the band edge to 12.8 kHz placed at the center of the band. The downlink is divided into 28 channels the narrowest is 3.2 kHz placed at the center of the band and the widest is 12.8 kHz at the center of the band. All channels uses GMSK modulation scheme similar to GSM. The downlink includes a spread-spectrum beacon signal used for fast device frequency and timing acquisition. It also carries information that enables downlink multicast service.

The CSS mode allows a frequency reuse of 1. It features variable spreading factors from 32 to 4096 with a chip rate of 125 Ksps. This mode provides positioning capability by locating uplink transmissions received by multiple BTS using time difference of arrival (TDOA) techniques with 10 – 100 m accuracy.

The LoRa narrowband mode provides for over 160 dB maximum coupling loss whereas the CSS mode provides lower MCL that tops at 160 dB.

4.2 SigFox Cooperative Ultra Narrowband (C-UNB)

SigFox technology in the uplink is based on ultra-narrowband channels of 160 Hz called ad-hoc micro-channels. There are 1250 such micro-channels in 200 kHz bandwidth where each micro-channel has a pseudo-random center frequency in the full 200 kHz band. Each micro-channel is modulated with D-BPSK to leverage existing sub GHz radio chipset market for low cost devices. The data rate per micro-channel is 160 bps. In the downlink, the subchannel bandwidth is 600 Hz channels with bit rate of 600 bps using 2GFSK modulation scheme. C-UNB is primarily an uplink technology as the MAC PDU support between 7 – 25 bytes in the uplink and 1 – 8 bytes in the downlink.

The device randomly selects three uplink micro-channels and transmits three repetitions of the data to increase robustness. C-UNB does not support device attachment to any base station and the device transmits without knowing which base station is in its range. All base stations listen to the same 200 kHz band. This allows for cooperative reception by multiple base stations where a message sent by a device is received by one or more base stations.

Transmission in the downlink is based on ‘time-delayed piggy-backing’ where downlink packets are stored in the core network and forwarded to the device after an uplink transmission. C-UNB does not support a paging mechanism and there are no means to wake up a device to push downlink packets towards it. In the case of multiple receptions by several base stations, the core

network selects the most appropriate base station for transmitting the downlink packet. There is no MAC-level acknowledgement in C-UNB which is left for applications to implement and manage through the application server.

C-UNB provides about 164 dB maximum coupling loss in both uplink and downlink with 24 dBm and 34 dBm transmit power, respectively.

4.3 Huawei/Neul

Neul has been developing its own IoT access protocol called Weightless which targeted TV whitespace bands in its broadband version and ISM band in its narrowband version. After the acquisition by Huawei, Neul proposed to GERAN a narrowband technology to slot into existing GSM channel allocations as well as potentially into LTE guard bands that are created by the null sub-carriers.

The uplink physical layer consists of 36 uplink sub-channels of 5 kHz for total channel bandwidth of 180 kHz. Each sub-channel is individually modulated with D-QPSK, D-BPSK or GMSK. Uplink sub-channels can be bonded by x2, x4 or x8 sub-channels and are used in a similar manner to OFDMA technology. The maximum data rate for a bonded sub-channel is 45 kbps (minimum per channel is 250 bps).

In the downlink, each 180 kHz channel is divided into 12 downlink sub-channels spaced by 15 kHz. Each sub-channel is individually modulated with BPSK, QPSK or 16QAM. The maximum data rate per sub-channel is 36 kbps for a total downlink rate of 432 kbps (minimum data rate per sub-channel is 375 bps). One downlink channel is reserved for synch /broadcast for network acquisition.

Qualcomm – NB-OFDMA

4.4

This access technology features narrowband OFDMA in the downlink and SC-FDMA in the uplink similar in many ways to LTE. The 200 kHz channel is divided into 72 active sub-carriers of 2.5 kHz in bandwidth with 10 kHz guard band at either end of the channel. This results in relatively long symbol size, where a single NB-OFDMA symbol is as long as 6 LTE symbols. In the time domain, the frame length is 1 second which is divided into a number of slots. The downlink includes a total of 171 slots (163 normal which carry data and 8 special slots for synchronization). The uplink includes two frame structures: structure 1 for normal cells with radii less than 8 km and structure 2 for large cells with radii up to 35 km. Uplink frame structure 1 consists of 142 normal slots and 24 extended slots where as frame structure 2 consists of 137 normal slots and 24 extended slots. NB-OFDMA allows for sub-carrier hopping to average inter-cell interference and to allow frequency reuse one deployment where all sub-carriers are used in every cell.

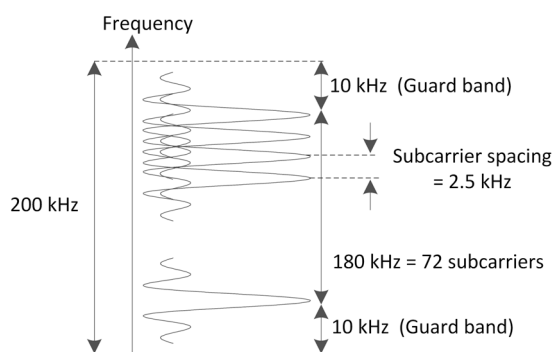


Figure 3: NB-OFDMA downlink and uplink frequency domain structure.
[Source: Qualcomm]

NB-OFDMA provides about 164 dB of maximum coupling loss with BPSK modulation to the cell edge.

5 The Implications

Wide-area IoT access technologies approach device connectivity from opposing directions. From one direction, LTE-M strips out many of the features required for mobile broadband connectivity to reduce cost and better match IoT application requirements especially for stationary devices. For example, LTE-M reduces channel bandwidth, defines single antenna operation, modifies medium access control layer to meet the intermittent, low data rate characteristics of many IoT applications. However, many of the fundamental design aspects of LTE cannot change which limits the extent to which LTE can be adapted to meet IoT application requirements.

From another direction, LPWA technologies are designed from the start to cater to IoT applications with an optimized air interface. LPWA are optimized for intermittent low-data rate transmissions. The access protocol is designed to support a large number of devices without coordination from the base station (or gateway) and build redundancy in transmissions to increase the robustness and reliability of the link. The access to the air interface is not scheduled, but rather based on contention which is typical of many systems operating in license-exempt spectrum. LPWA technologies build a higher system gain than today's GSM and LTE systems for longer reach, a feature that the evolution of LTE for machine communications is working to address.

In the balance, there are tradeoffs between these technological approaches that can only be viewed within a larger context that is not limited to the air interface. Some considerations include the following:

The network 'backend'. This is a general term we use here to denote functions such as network control and management, device management, billing, security, and other such functions that are required for both operational and business processes. These are critical functions that have been in development for many years by service providers and are optimized for consumer services. Adapting these functions for IoT applications carries both advantages and disadvantages for established mobile operators. On the other hand, LPWA systems are relatively new and the network backend remains fragmented and does not measure to the same level of maturity as that of the mobile network. However, there is no burden of legacy which provides an opportunity to define optimized systems and solutions in this space.

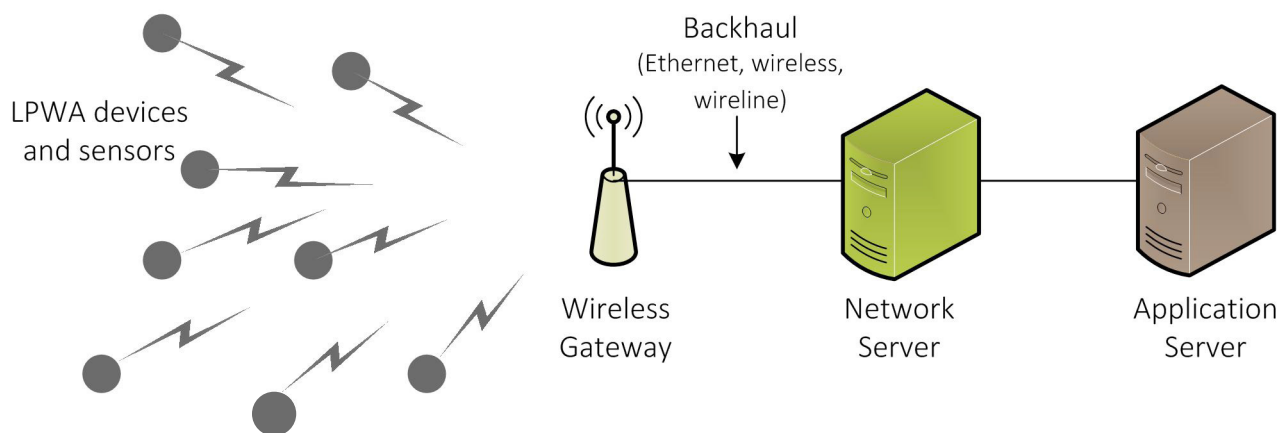


Figure 4: LPWA networks architecture.

We foresee a significant degree of innovation related to the LPWA network core/backend coming to market over the next 2-3 years. This is primarily due to 2 reasons: First, the Greenfield nature of LPWA networks provides an opportunity to design solutions taking full advantage of cloud services delivery models, data management architectures and intelligent data processing technologies; and second, the relative decoupling from the 3GPP protocols/standards that have led to very specific product and solutions architectures in the core of the network.

With legacy constraints relaxed, the new core/backend solutions will emphasize agility, costs elasticity and scaling efficiency, which in turn will allow the delivery of IoT-centric services with superior cost/value economics. They will also tackle the challenges around IoT service security and synergistically integrate the network into the value chain of different industry verticals.

This will be a space to watch closely especially as LPWA technologies are set to bifurcate as they are brought under the 3GPP umbrella to accommodate mobile network operators. As LPWA solutions converge towards industry standards, the resulting core is likely to be different from the solutions deployed today. It is this combination of alternative wireless access technologies, as is the case with LPWA, along with fundamentally different core/backend systems, that would enable the business case for the deployment of certain IoT services.

Spectrum. Sub 1 GHz licensed spectrum is expensive and owned by mobile network operators. 2G technologies typically operate in older grants of this spectrum while newer grants represented in digital dividend spectrum typically operate LTE. Whatever the case, operators around the world plan to reform this spectrum to LTE eventually as 2G and 3G technologies near their end-of-life cycle (for example, in the United States, AT&T will turn off 2G while European operators will tend to turn off 3G first). Hence, narrowband technologies will have to coexist with LTE in a defined spectrum or operate in unlicensed spectrum such as the ISM band. MNOs have based their business model on service reliability and high availability would seek to deploy IoT solutions in licensed spectrum bands as there's always the risk that interference in license-exempt spectrum would reduce reliability and service availability. This is bound to raise the cost of service. On the other hand, LPWA technologies are designed to deal with interference by defining an air interference with greater tolerance, redundancy and robustness than cellular technologies as it supports low data rate. These two approaches are bound to collide although they can be viewed as complementary whereby applications with intermittent low data rate can use license exempt spectrum leaving applications requiring frequent access with service assurance to use licensed spectrum.

The business case. A critical challenge in enabling IoT service has been validating the return on investment. Assessing the costs and benefits of IoT is a challenge due to many reasons that transcend the cost of the module which has been the focus on the telecom industry. Enabling IoT requires integrating connectivity to derive intelligence from which value is extracted. Connectivity is fundamental but it is not the sole driver for adoption. Yet, connectivity introduces both capital (system integration) and operational expenditures that must be accounted for by the user. The cost of connectivity is then a key hurdle that must be cleared. The lower the cost of connectivity, the fewer objections or hurdles IoT would face.

While a comprehensive overview of the business case is beyond the scope of this whitepaper, we touch upon the cost of the device which, as stated, has been a focus for the industry. The general requirement for narrowband technologies as specified by GERAN and 3GPP is below \$5/module.

Several LPWA system and module providers claim meeting this number and reaching values as low as \$1 in large volumes. This is a drastically different from the cost of cellular devices where as we mentioned a GSM/GPRS module costs around \$10 and an LTE module close to \$50.




Figure 5: Cost structure for IoT devices.

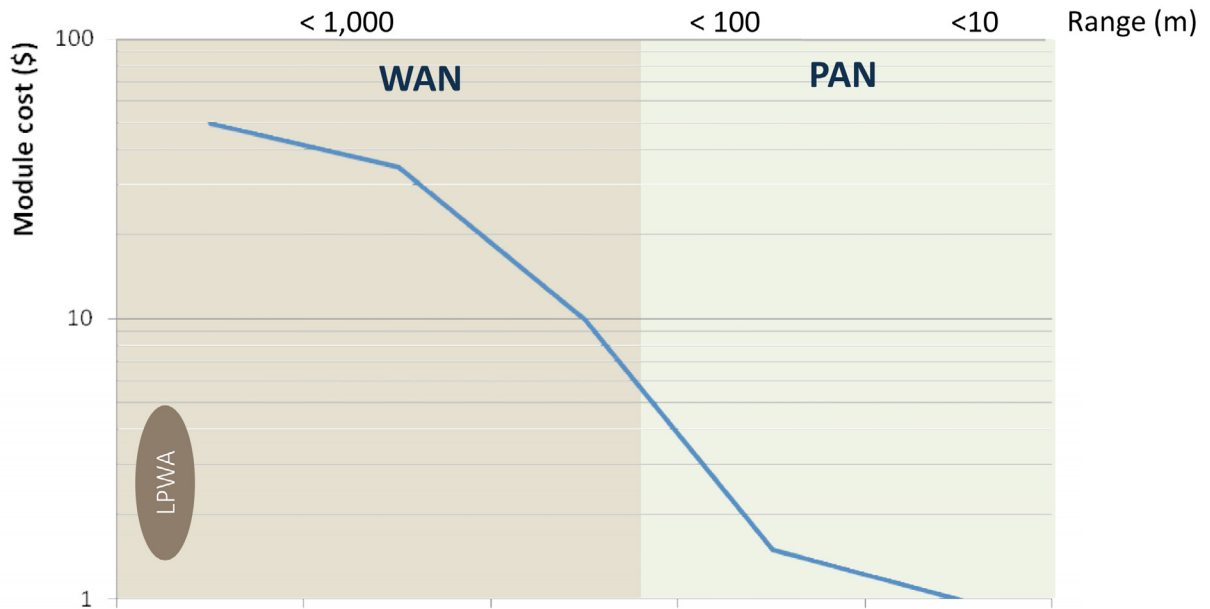


Figure 6: Device cost in IoT applications.

Mobile network operators rely on an existing framework for providing service while LPWA challenge this framework with new operational and business models. While legacy systems provide an advantage in the short term, they fail to meet long-term objectives. This is where the opportunity for LPWA lies provided it can prove a positive business case and sufficient operating performance. MNOs that would have the capability to deploy LTE-M will need to carefully weigh

their options as their cost structure may exceed the threshold required to enable some IoT applications, especially ones based on very low and intermittent data rates. On the other hand, LPWA operators would need to ensure that the business model and cost of service will lead to profitability.

Table 2 Comparative assessment of wide-area IoT technologies.

	Advantage	Disadvantage
LTE-M evolution	<ul style="list-style-type: none"> • Existing ecosystem of operators • Ability to leverage existing LTE network operation processes and framework for core network (upgrade still be required) • Licensed spectrum • Higher throughput performance • Reliability and service level agreements • Established infrastructure 	<ul style="list-style-type: none"> • High cost base (capex & opex) • Short range • High power consumption (in relationship to narrowband technologies)
Narrowband technologies / LPWA	<ul style="list-style-type: none"> • Designed for IoT device connectivity requirements: <ul style="list-style-type: none"> - High system gain for long range and fewer sites - Efficient medium access control layer - Efficient power management for long field operation on batteries • Business models and pricing schemes aligned with IoT business case requirements • Low cost of devices and service • Scalability to support high number of devices 	<ul style="list-style-type: none"> • Nascent and evolving ecosystem • Fragmentation: many technologies vying for prominence • Spectrum: license-exempt spectrum raises questions on reliability of service • Unproven: LPWA has few deployments today. Scalability, business model, and many other factors remain to be validated

6 Concluding Remarks

Wide-area IoT connectivity is on the cusp of a major shakeup that will unfold in the next few years. The shortcomings of today's cellular technologies are evident with the limited proliferation of wide-area IoT and the potential opportunities that new technologies can unleash. IoT services are fundamentally different from consumer broadband services. Yet, the wireless industry has primarily worked at retrofitting existing network and service models designed for consumer broadband to running M2M/IoT networks with limited success to date. Narrowband or LPWA technologies are designed from the ground up to cater to low-power, low-data rate, and longevity in the field. They are also designed for high scale and long range to enable a better business case in comparison with existing cellular technologies. LPWA powered by new core/backend technologies provide a new way for delivering services that is better optimized to the application requirements. However, cellular technologies have key strength in an established and vibrant ecosystem, licensed spectrum, and an infrastructure on which to build and evolve which the LPWA ecosystem is working to create. Cellular technologies are advancing to support device communications along their own roadmap. These trends are creating interesting dynamics as the boundaries for collaboration and competition are being defined with high stakes to decide the winners for a market valued in the trillions of dollars.

7 Acronyms

2G	Second generation
3G	Third generation
3GPP	Third Generation Partnership Project
4G	Fourth generation
AAA	Authentication, Authorization, and Accounting
BPSK	Binary phase shift keying
Cat	Category
CSS	Chirp spread spectrum
C-UNB	Cooperative Ultra Narrowband
D-BPSK	Differential binary phase shift keying
D-QPSK	Differential quadrature phase shift keying
DRX	Extended discontinuous reception
eUICC	embedded Universal Integrated Circuit Card
FDD	Frequency division duplex
GERAN	GSM EDGE radio access network
GFSK	Gaussian frequency shift keying
GMSK	Gaussian minimum shift keying
GSM	Global System for Mobiles
IoT	Internet of Things
ISM	Industrial Scientific Medical
IWF	Interworking function
LAPI	Low access priority indicator
LoRa	Long Range
LPWA	Low power wide area
LTE	Long Term Evolution
LTE-M	LTE Machine
M2M	Machine to machine
MAC	Medium access control
MNO	Mobile network operator
MSISDN	Mobile Station Integrated Services Digital Network
NB-OFDMA	Narrow-band OFDMA
OFDMA	Orthogonal frequency division multiple access
PDU	Packet data unit
PSM	Enhanced power saving mode
QAM	Quadrature amplitude modulation
QPSK	Quadrature phase shift keying
Rel	Release
Rx	Receiver
SC-FDMA	Single carrier frequency division multiple access
TAU	Reduced tracking area updates
TDOA	Time Difference of Arrival

XONNA PARTNERS

Internet of Things

The Turning Wheels of IoT Investments

Dr. Riad Hartani, Frank Rayal,
Dr. Dean Sirovica

February 26th, 2015

Investments are steadily flowing into the Internet of Things (IoT) ecosystem as investors attempt to match market analysts' optimism about the sector. As an example, French-based SigFox has recently raised \$115 million (€102m) from Telefónica, NTT Docomo, and SK Telecom in addition to other investors. SigFox developed a low-power, long-range technology to connect devices and has been rolling out networks to provide connectivity infrastructure services for different industrial and commercial applications such as utility smart meters among other applications. This technology, which is commonly referred to as Low Power Wide Area (LPWA) connectivity, is considered as complementary to what mobile network operators provide in their current M2M offering which is skewed for high data-rate applications that justify the higher charges MNOs want to bill for this service. This investment follows another made a few months earlier when Huawei acquired Neul, a UK-based company developing a LPWA protocol.

The optimistic investor sentiment has permeated all aspects of the complex IoT space which spans devices and connectivity technologies, applications, platforms and services. Corporations view IoT as the next phase of growth where new business opportunities will be created similar to what happened in the Internet space connecting people. On the other hand, the dynamic environment of the IoT space provides genuine opportunities for startups to make their mark and to profoundly impact the status quo where IoT applications have often stuttered because of poor business case, a complex ecosystem and complex processes that did not allow all elements of the ecosystem to derive value. Yet, there are two areas of real growth seen today: one related to industrial IoT where utility smart meters have been leading and another related to wearable technologies that leverage the smartphone as a connectivity gateway.

From a strategy and investment perspective, corporations are jockeying to secure and solidify their home turf against competitors and to establish new opportunities for growth. They are investing and partnering to accelerate product development, expand service offering, license technology, and acquire knowhow. An example is Qualcomm's acquisition of CSR for about \$2.5 billion to secure leadership in Bluetooth for short-range connectivity. Qualcomm had prior made investments in other chip companies such as startup Ineda which is developing an ultra-low power system on chip (SoC) for wearable devices. GE on the other hand partnered with and invested \$30m in Quirky to develop connected home solutions and services while Samsung acquired SmartThings to gain a platform for connected devices. Telecom operators have also been on the lookout for acquisitions in the IoT space as a way to get into adjacent markets via non-organic growth. The automotive sector has been a primary focus. Most relevant examples include Vodafone's acquisition of Cobra Automotive Solutions and Verizon acquisition of Hughes Telematics.

Corporations are setting up funds for IoT investments as well as investing in incubators of IoT focused startups. Samsung setup a \$100m accelerator fund for IoT investments into startups in the \$100k to \$2m range. Cisco has allocated at least \$250m for IoT startups in addition to other investments into accelerators and private equity funds focused on IoT. In particular, Cisco has focused on cyber security with a number of acquisitions to solidify its position in an area critical to the take-up of IoT services. For example, Cisco invested \$2.7bn in SourceFire and made additional investments in ThreadGrid, Cognitive Security, and Virtuata. Intel on the other hand opened an IoT lab (Ignition Lab) in Swindon, UK, in June 2014 to focus on smart cities including applications for buildings, retail and transportation. It also acquired Basis Science which makes wristband health trackers for more than \$100m.

Telecom infrastructure vendors have been ramping up their in-house development of IoT solutions as well as making strategic acquisitions. The acquisitions have focused on areas around cloud platforms and network OSS and analytics solutions for IoT applications. The objective behind these acquisitions is to optimize the solution offering to cater to a new wave of IoT deployment models. Examples would include Ericsson's acquisition of MetraTec which is related to OSS and Cisco's acquisition of Tail-f which is related to infrastructure virtualization and optimization. We foresee the next wave of investments and acquisitions to include IoT specific solutions in different business verticals, along with areas related to infrastructure development.

Chipset and subsystem IoT providers are faced with the most significant challenges in terms of where to focus investments given the fragmented nature of the IoT connectivity market. The leading vendors are likely to monitor and hedge the market via their venture capital investment arms. Intel Capital would be a primary example, having been the number 1 venture capital investor in IoT in 2013. Qualcomm setting up a China-centric IoT investment fund with \$150m would be another example.

Venture Capitalists have also been moving into the IoT with over \$1.1 billion in funding in 2013 – a 57% increase over 2012. More than \$1.4bn has been pumped into wearables since 2009, of which over \$500m was invested in 2014 alone. Some venture capital firms have had IoT investments as a priority over the last few years, and have been leading the latest rounds in this space. This includes Intel Capital, True Ventures, Qualcomm Ventures, Cisco Investments and KPCB as some of the most active investors in 2013 and 2014. Health and wellness, location services, and healthcare are the highest investment sectors garnering over 50% of total VC investments. In the last year there has been a notable increase in Angel investor deal flow for startups related to IoT. Activities are relatively high at the seed stage with an almost even distribution between series A, B, and C. There is also a relatively high percentage of strategic investors reflecting the need of the IoT market for large enterprises creating a market for IoT related products and services. Companies such as Jawbone, Fitbit, and mc10 are among the highest funding recipients. Another area that has seen significant VC and corporate interest is the platform space which is where many consider the value of IoT will reside. This space has been vibrant with many entrants as well as corporate investments to have a lock on a critical piece of the IoT value chain.

Machine-to-machine mobile virtual network operators (MVNOs) today account for 4% of all MVNOs. The development of IoT will progressively lead to strong growth of IoT-centric MVNOs, targeting specific industry verticals. This is a further evolution of the data MVNO model. Various startup MVNO operators are in early launch stages while the large Internet players (Google, Ali Baba, Amazon, etc.) are at various stages of validation of these IoT centric MVNO technologies and business models.

The outlook for investment environment in IoT continues to be promising in 2015 as several of the trends that emerged in 2014 will further drive interest and value in this wide ecosystem. Some of these trends include the alliances established to facilitate interoperability between IoT devices (e.g. the AllSeen Alliance established in December 2013 and the Open Internet Consortium established in June 2014), as well as new technologies that will emerge for a number of other projects being standard-based or resulting from independent development.

We conclude this review of IoT investments landscape by pointing to a few developments to watch for in the 2015 timeframe and beyond:

- Telecom operators who have been betting on 3GPP-centric technologies as the main wide area IoT connectivity model are compelled to revisit their assumptions and develop LPWA-related strategies. The specific nature of these wireless network deployments and the amount of capital required for such strategies would require the participation of large corporate venture funds from various industrial conglomerates as well as telecom operators' own investment participation.
- IoT services provide the opportunity for new MVNO models to develop, led by startups, leading industry vertical players, or large Internet players. Significant investment is likely to be made in this space over the next 2-3 years.
- Connectivity in the local area is likely to remain fragmented with various technologies in use, and as such, major chipset and systems vendors are likely to pursue a multi-technology strategy, and monitor the market via their venture capital arms investing in IoT startups globally.
- The IoT platform business will continue to evolve but would remain fragmented for some time with the large Internet and infrastructure players aiming to dominate it in select industry verticals via targeted acquisitions.

Read the original article on: <http://www.tmtfinance.com/content/internet-things-part-2-turning-wheels-iot-investments-xona-partners#ixzz3TA7mkc88>

XONA PARTNERS

Internet of Things

Coming of Age

Dr. Riad Hartani, Frank Royal,
Ananda Sen Gupta (Xona Partners),
and Ricardo Tavares (TechPolis)

January 26th, 2015

Preamble

The Internet of Things (IoT) is by definition a vast topic that encompasses multiple markets, technologies, and disciplines. It is impossible to do service to this field in a single paper, which makes our attempt to accurately characterize the market and call out important issues in this short paper particularly ambitious. Nevertheless, we lunge forward with an overview of some of our thoughts and observations while we admittedly leave many areas uncovered.

Connecting devices and ‘things’ to the Internet is a natural evolutionary step after two decades of focusing on connecting humans to the Internet. However, while the term ‘Internet of Things’ may go back to 1999, elements of IoT preceded that date. In its earlier form, the focus of IoT was on sensors and tracking devices – an example can be found in fleet tracking technologies using GPS such as Omnitrac. Early applications focused on commercial, industrial and even military sectors, illustrating the difficulty of narrowing down a precise definition of IoT’s scope.

The current term for IoT is in fact, much broader than the original. The recent growth in connected consumer devices, which include wearables and connected home, health and car applications, has skewed the definition of IoT towards the consumer sector. The impetus for this change is due to the proliferation of smartphones and data services that provide connectivity to remotely controlled devices that transfer rich multimedia content. The developments in wide-area connectivity are mirrored by equally important evolution in highly scalable compute platforms for low-cost storage and data processing capabilities, which also plays a fundamental part in propelling data science applications to provide value added services that improve the business case of IoT. This is a critical point, as many IoT applications would fail on a commercial basis without the additional value derived from services that are enabled through the cloud.

Together with the promise of IoT comes a series of obstacles that combined to slow down the rate of adoption of many smart technologies. IoT applications are broad, fragmented and siloed in specific verticals where multiple competing technologies vie for prominence resulting in incompatibility. The topics of security and privacy become complex, and often requiring intervention to frame a regulatory context that provides direction for further development. From this perspective, IoT is an evolutionary process that will exhibit varying adoption rate in each silo while the market works its way through the challenges.

In this paper, we layout an ecosystem reference model for IoT and provide a brief overview of some key challenges and evolving trends that characterize each layer.

The IoT Ecosystem

To conceptually define IoT, we present a five-layer functional model that includes devices, connectivity, applications, platforms, and services (Figure 1).

Devices: Sensors, identifiers and gateways are types of IoT devices used to collect and convey information. Devices are designed and deployed to meet the application use case requirements. They can range from simple identifiers that provide specific information on the object, to complex devices that have the ability to measure (sensors) and process data (gateways). The application, use case and deployment scenario places requirements on the device such as size, weight, power consumption, life of operation or deployment. This in turn impacts the connectivity of the device

to the network. A variety of IoT devices have emerged in various business verticals, starting in the utility / energy sector to include today devices in the health, transportation, home and finance ecosystems among others.

Connectivity: Devices can be connected directly to the network, or indirectly through another similar device (mesh) or a gateway that is provisioned to support multiple devices. Connectivity can be through a number of physical media such as copper, fiber optical cable or over the air through a number of wireless technologies. One of the challenges in IoT is the proliferation of connectivity standards, which is a common symptom of the breadth and fragmentation of IoT application requirements. These standards span the entire logical protocol stack through layers 1 – 7. Examples of connectivity would include the traditional 2.5/3/4G networks, as well as various local area solutions (Zigbee, Wi-Fi, Bluetooth, others) and low power wide area solutions (e.g. Weightless) among others.

Applications: Applications define the use case of the device and include all the necessary functions required to make use of the device for the intended purpose including the hardware and software architectures. IoT application stores are emerging with applicability to specific industry verticals, with the health wearable devices being a recent example.




Figure 1. IoT ecosystem reference model.

Platforms: devices and connectivity requires a platform to provide a service. Platforms are used to provision devices, manage and control them. They are used for billing and fraud detection. Platforms also provide the means to customize functions and data according to the requirements of end users. From this perspective, platforms allow the IoT infrastructure to perform as required.

Services: This references the IoT service to the end-customer. The service provider leverages all the downstream elements in this value chain: platforms, applications, connectivity and devices. The service provider can be the same or different from the platform and application provider. Examples would include automotive automated diagnostic, medical geriatrics and remote power consumption optimization.

The IoT Connectivity Model

To help drive conclusions and observations on IoT development, we intersect the IoT reference model presented above with a model for data flow, which can simply be modeled by three stages: data creation, transmission, and consumption.

Data creation: Data is generated by different types of devices, as described earlier. Data has specific characteristics such as rate, volume, latency, and frequency. For example, video surveillance has high data rate whereas SCADA systems have low bit rate. Taking this example further, we note that in many SCADA applications, the latency has to be very low to accommodate specific requirements of an application such as a fault in an electric transformer that require instantaneous switching of electric current to avoid damages while there is higher tolerance to latency in video applications.

Data transmission: The characteristics data place requirements on transmission in terms of bandwidth, latency, compression, encoding, multiplexing, privacy and security. Thus, different types of pipes are used for transmission as outlined above in connectivity: GPRS, 3G, 4G, LPWA, IP, P2P, DSL, satellite, fiber, etc.

Data consumption: Data is consumed by different segments of end users according to the application. This can be through simple systems that involve the user directly interacting with device, for example, interacting with a wearable through an application on a mobile device or tablet. Alternatively, sophisticated techniques based on data sciences can be used to derive additional information, which can be used to the mutual benefit of the end user and a third party. A homeowner may install a Google Nest thermostat, which she can control remotely; however, the data can also be shared with the utility company to control temperature within certain bounds during peak hours.

The intersection between the IoT and data reference models is used to develop a number of observations and conclusions on the state of the IoT market as we outline below.

Observations on IoT Market Space

We can deploy the conceptual IoT framework above to model developments across the ecosystem layers starting with devices and connectivity and working upwards towards platforms and services.

To start, we note that the IoT use case determines requirements for devices and connectivity. Device characteristics such as size, weight, placement, mobility, power and communication characteristics as defined by the application drive the connectivity requirements. The great variety in use cases in each vertical market (for example, automotive, home, health, industry, etc.) has resulted in proliferation of connectivity standards.

1- Proliferation of connectivity standards: Connectivity standards can be divided into different categories depending on fundamental characteristics. In our model, we used the following three categories: Spectrum requirements (for wireless connectivity; devices can be connected through wireline technologies such as PLC); and range, power and cost which are highly correlated. 3GPP standards such as GPRS, UMTS and LTE are licensed band access schemes that rely on high power for long range, consequently are relatively expensive in comparison with other connectivity techniques. On the other hand, technologies such as Bluetooth are meant for short-range communications in unlicensed spectrum and are low on power consumption. Various LPWA proprietary solutions have also recently emerged, mostly in unlicensed sub-1 GHz spectrum but also in some licensed bands. Wi-Fi relies on higher power and provides longer range than Bluetooth albeit at a higher cost.

In recent years, advancements in silicon technologies such as 28 and 14 nm processes have significantly reduced power consumption to allow ever-smaller devices with less battery requirements to come to market. Coupled with the maturity of smartphones, this led to a great jump in interest in wearables and personal connected devices.

2- Commoditization of devices: Devices and connectivity continue to march on a downward slope of cost reduction (Figure 2). This is essential to enable the business case for IoT applications. The challenge to device manufacturers is how to differentiate from competition. Our observation in this space is that software applications and platforms, including operating systems, are the essential leverages used by device manufacturers to differentiate (e.g. Apple/iOS, Google/Android; Samsung attempt at differentiating through Tizen, and in a similar way with Alibaba and Xiami's own platforms design).

	Xiaomi Mi Band	Fitbit Flex	Jawbone Up
Steps taken	✓	✓	✓
Calories burned	✓	✓	✓
Distance traveled	✓	✓	✓
Active time	✓	✓	✓
Sleep time	✓	✓	✓
Sleep quality	✓	✓	✓
Map routes		✓	
Average pace		✓	
3rd party apps		✓	✓
Diet tracker			✓
Waterproof	✓	✓	✓
Bluetooth	✓	✓	
Social sharing	✓	✓	✓
Alarm	✓	✓	✓
Notifications			✓
Indicator lights	✓	✓	✓
Price	\$13	\$99	\$79

Figure 2. Device commoditization.

3- Commoditization of connectivity: Low-cost connectivity is essential to enable the business case of most applications. There are many variants of connectivity including wireline and wireless technologies. The lowest cost wireless connectivity leverages license-exempt spectrum over short distance (Figure 3). Wearables, for example, leverage Bluetooth to connect with smartphones. Alternatively, some consumer devices rely on longer-range license-exempt technologies such as Wi-Fi for greater range. Central hubs for connectivity and routing are deployed to tether over longer distances for remote control and monitoring. Where mobility is required, wireless technologies in licensed spectrum can be implemented albeit at a higher cost.




Figure 3. Commoditization of wireless technologies.

4- Emergence of long-range low power wireless technologies: We see an opportunity for very long range wireless technologies that are low power, low cost and work over long range (Figure 4). Such technologies are now on the market but are yet to prove their commercial viability (for example, Neul Weightless, SigFox Ultra Narrow Band, Semtech LoRa, and On-Ramp). These technologies often assume the buildout of a parallel IoT network to the mobile network. The IoT network is operated as a private network on a subscription model of per device/message basis for low fixed cost pricing.

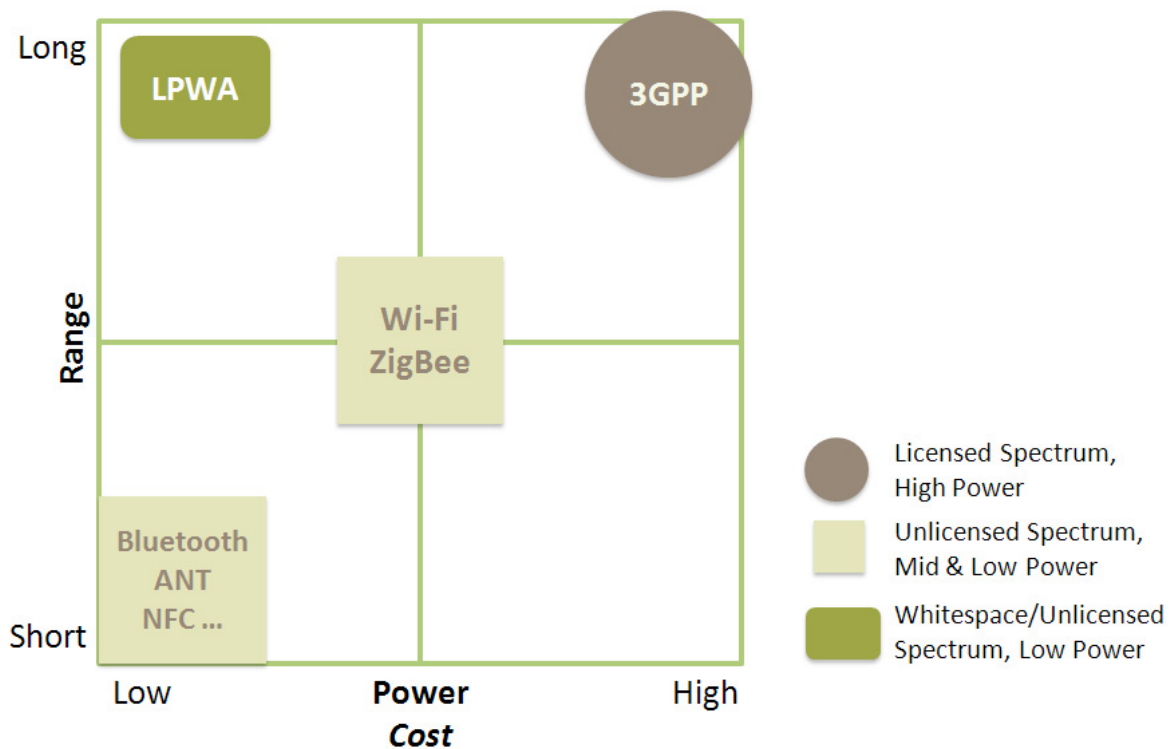


Figure 4. IoT Wireless connectivity.

5- Competition and harmonization of connectivity standards: Connectivity standards have been progressing slowly but steadily. The challenge is not really in the definition of these standards, but more in terms of the number of variety of competing and complementary standards, as well as the conflicting interests of the industrial groups behind the various standards. Although harmonization is ongoing, it is very likely that IoT solutions will face challenges for rapid mass adoption. The development of interworking platforms with open APIs will help alleviate some of these challenges by allowing interoperability of different standards or different implementations of the same standards. This is not only the case for physical and link layer standards, but also includes aspects related to applications and services running on top of the IoT ecosystem.

6- Partnerships and alliances to win the IoT platform war: The development of IoT solutions is inherently about the development of ecosystems around offered solutions. Such ecosystems are built via tight and loose partnerships between the various industry players. The leading players will aim at controlling the ecosystem by providing a platform that would host IoT applications, and over which IoT services will be built (Figure 5). As in any platform model, such as those in smartphones and the Internet, the key is to increase the adoption of the platform. Various models are being put in place to achieve this, via the development of open source IoT connectivity and interworking software, open APIs to plug into the platforms, and SDKs to develop services on top of the platform. We foresee the emergence of fragmented alliances over the next few years, across industry verticals, with a focus on advancing specific IoT platforms, but progressively evolving towards selecting winners, as it's traditionally the case for Internet-centric business models. Various contenders are already in the game to achieve this, including the Internet platform players (Google, Apple, Amazon, etc), the lead industrial players with a specific

vertical focus (e.g. GE for industrial Internet), and to some extent certain mobile operators with a strategy towards Internet-scale OTT deployment.

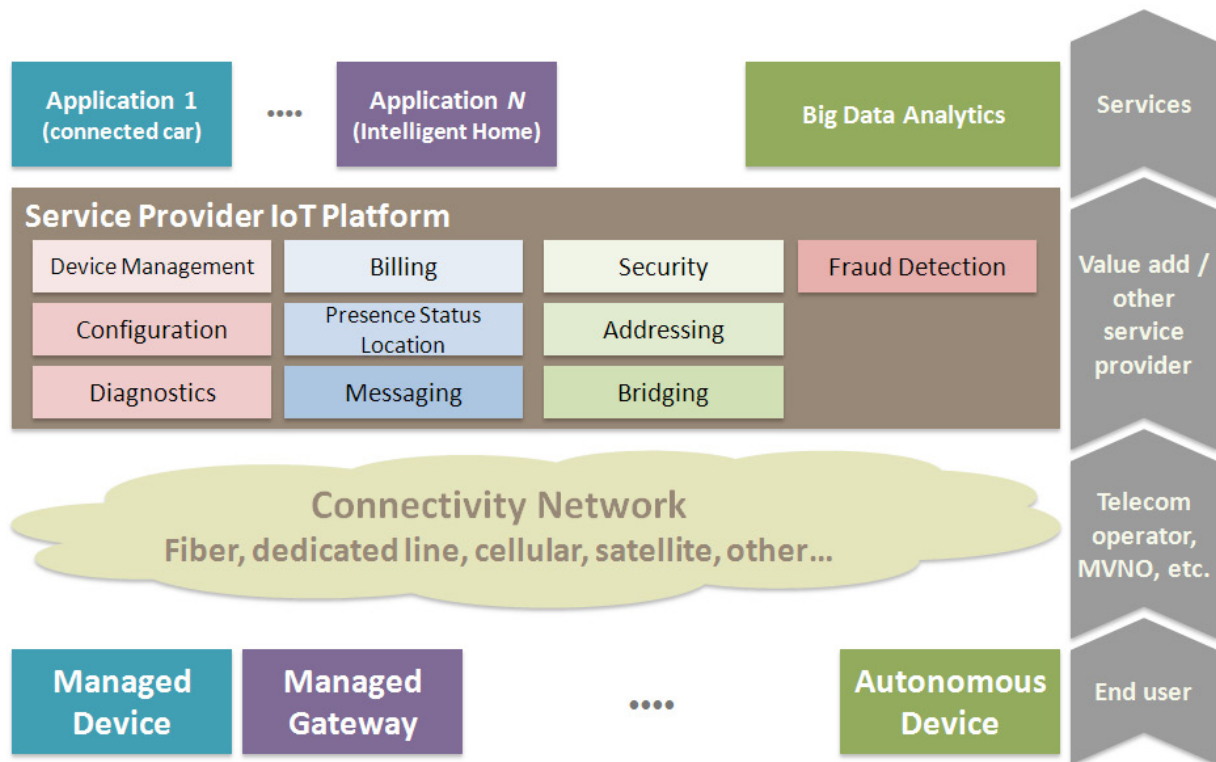


Figure 5. Value appropriation through platforms.

7- Emergence of new MNO and MVNO service models: A key dynamic of the IoT market that needs to be highlighted is that the majority of 'value' in any IoT application lies not in the simple carriage of data, but in the provision of an overall service. For example, a wide-area wireless enabled home security system represents a significant revenue opportunity for a mobile or virtual mobile operator, including revenues from device sale, installation, and monthly service fees. However, the data traffic revenue that such a solution generates is likely to be relatively small in comparison. In a similar fashion, a connected health solution would include the connectivity network as well as the platform to manage the solution, interfacing with the various stakeholders in the health solution value chain. The story is the same for many other IoT applications: the real opportunity for mobile operators lie in moving up the value stack and away from the simple provision of data carriage services. Basically, to provide IoT service, the ecosystem will be complex, multi-party and heterogeneous in nature. The mobile network operator has to provide the connectivity and IoT management for value added. IoT solution provider (either over-the-top IoT service provider or mobile service provider who offers IoT solutions) has to integrate all the components of the ecosystem for the end-to-end IoT solution.

8- Extracting value through data sciences: As businesses evolve to leverage the huge amounts of data assembled – mining and learning through such data as well as optimizing communication between those producing it and those using it brings significant opportunities around IoT business models. As such, the desired goal is to create a solid foundation architecture that is able to provide these optimal functional capabilities and a platform to overlay data science applications. This would include the various layers in the data value chain – optimized processing

through an acceleration of migrations to the cloud, scalable data management leveraging big data models and the use of customized data sciences solutions for business intelligence creation. This is complemented by a fundamental re-architecture of IT models within the businesses integrating IoT models. We are now witnessing the emergence of enhanced (and new in some cases) set of machine learning and data mining algorithms, specifically focused on clustering and predictive modeling in high dimensional spaces based on imprecise, uncertain and incomplete information, efficient statistical data summarization and features extraction algorithms as well as large-scale real-time data stream management. These tools will be at the core of the processing engines being commercialized or running in open source environment, and will aim, when applied to specific industry problems, at optimizing the existing business logic and augment it with new functionalities over time.

9- The policy and regulatory conundrum: We are witnessing two disparate trends in the policy and regulatory realm. On the one hand, an important new set of policies designed to deliver benefits to the public are encouraging IoT deployment. An example of this is the eCall schemes in the European Union and Russia. Their goal is to reduce deaths by preventing unnecessary car accidents and expediting assistance from the emergency services when accidents do occur. The schemes mandate connectivity between individual vehicles and other elements of the transportation vertical. The mandates are slow moving, but have great potential to accelerate “smart transport.” Another stream of policies enhancing IoT adoption is in the Smart City arena, where policy makers are introducing new systems for urban management in urban transportation (V2V, V2P, smart parking), disaster prevention and public security. On the other hand, IoT challenges some pre-existing regulations which require adapting, especially in the heavily-regulated mobile telecoms industry. Examples include numbering schemes, international roaming and SIM-card registration. While policies aiming to deliver public goods using IoT are accelerating its adoption, there are also a number of regulations slowing down or simply destroying emerging business models. The key is for policies to rely on general frameworks supported by private sector delivery, and for regulations to become more flexible to allow new business models to flourish. A positive sign has been the proliferation of regulatory agency consultations asking questions such as: Should regulators set technical standards if markets fail to do so? (Ofcom, UK). Does IoT require specific spectrum? (Arcep, France). Spectrum for IoT is likely to come to the fore during the debate on new frequency bands for 5G, which will catalyze this year’s World Radio Conference (WRC-15) in Geneva (November 2nd to 27th). Table 1 provides an overview of the institutional environment for IoT policy and regulation.

Table 1. Overview of institutional designs for IoT policy and regulation.

	Instrument	Agency	Issues	Examples
Policy	• Legislation	• Parliament	<ul style="list-style-type: none"> • Data privacy & security legislation • Connected Car mandates 	<ul style="list-style-type: none"> • EU Parliament Directive 95/46/EC on personal data flow • EU eCall • Brazil 's tax breaks for M2M connections
	• Executive Order	• Executive Branch	• Connected Car mandates	• Ordinances implementing mandates
Regulation	• Statutory Regulation	<ul style="list-style-type: none"> • Telecom Regulator • Data Protection Regulator (Australia, Canada, HK) • Verticals regulators/Ministries 	<ul style="list-style-type: none"> • Permanent roaming of foreign SIMs in connected cards • Penalties for data privacy violations, rules for protecting data • Norms for smart grid 	<ul style="list-style-type: none"> • Brazilian regulator Anatel trial periods ordinances • TRA-UAE consumer data protection policy
	• Self-Regulation	• Trade associations, industry coalitions	• Child pornography, Advertisements	• GSMA agreement to block child pornography from mobile networks
	• Privacy by Design	• Individual Companies	• Allow users to opt out easily or protect privacy	• Safari's "Clear History"
	• Co-Regulation	• Multi-stakeholder: Civil society, government, industry	• Internet Governance	• ICANN's management of DNS

10- Evolution to 5G: In terms of timing and mass market adoption of advanced IoT solutions, it is very likely that this will converge and overlap with the specification and rollout of the first 5G networks. It is then natural that 5G specifications will have to take into account IoT requirements, either directly or via the complementary technologies that would form the future mobile ecosystem (including evolutions of Wi-Fi, LPWA, Zigbee, etc). As such, the LTE roadmap will continue to evolve to include new features that represent a precursor to those in 5G. For example, LTE-MTC in Release 13 aims to reduce power consumption of LTE devices for IoT applications and achieve low cost points by eliminating some of the broadband features of LTE (Figure 6). On the core, backend and underlying IT infrastructure, a gradual move towards virtualization, specific functionality enablement in private/hybrid/public cloud environment, and integration of big data analysis frameworks into network data management, will start appearing. All of these aspects will in essence contribute to bringing advanced IoT solutions and IoT centric business models to markets.

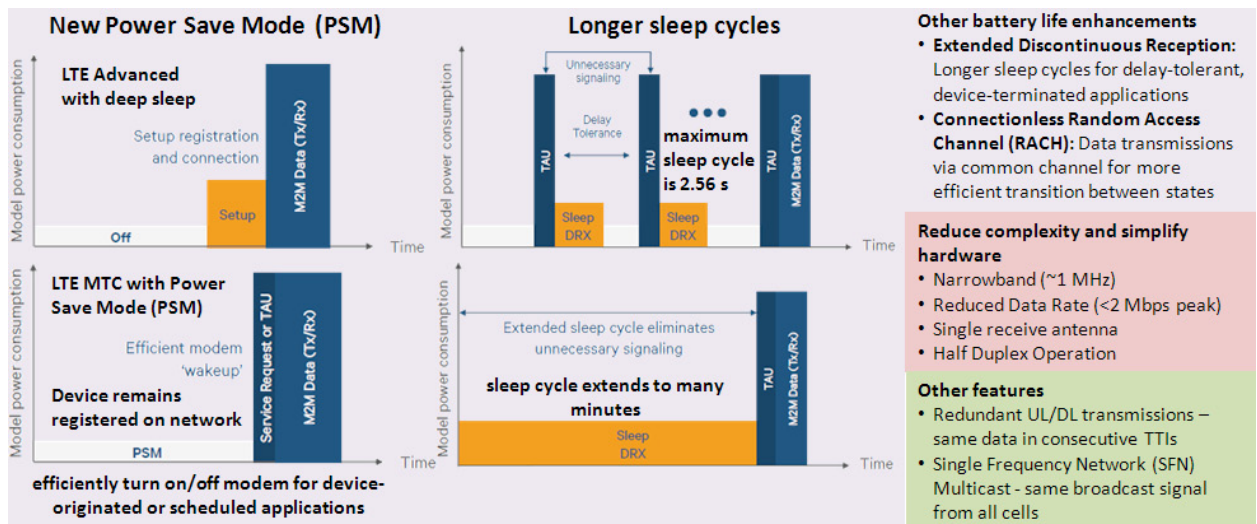


Figure 6. LTE-MTC features.

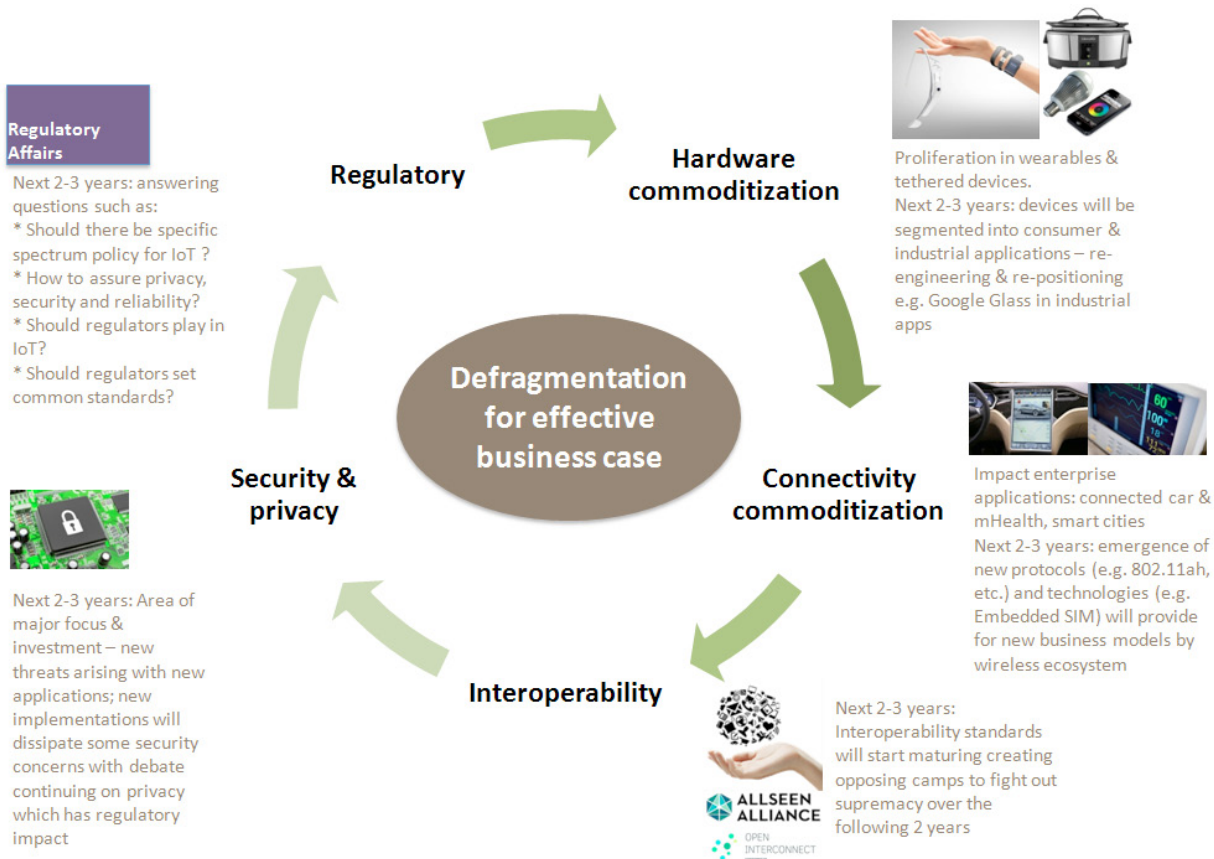


Figure 7. IoT evolution.

IoT – The Road Ahead

The IoT era has had various false starts, as far as mass adoption and progression to mainstream. The recent convergence of various trends including innovation in low power and low cost device technologies, scalable network connectivity as well as mainstream cloud and big data processing models, policies encouraging mass adoption in the transportation sector, have opened a new window for the emergence of IoT based value added services. In this paper, we took a systemic view of the IoT ecosystem which we divide into five layers and leveraged our experience with recent deployments of IoT solutions in select industry verticals, and working jointly with the various players in the IoT value chain, including device and chipset vendors, network connectivity providers, and suppliers of platforms for IoT service delivery to explore some of the most significant trends, both mid and long term, which we highlighted with implications on how the ecosystem will likely evolve and the underlying challenges and competitive positioning models that would emerge in this market.

Acronyms

3G	Third generation
3GPP	Third generation partnership project
4G	Fourth generation
5G	Fifth generation
API	Application program interface
DSL	Digital subscriber line
GPRS	General packet radio service
GPS	Global positioning system
GSM	Global System for Mobile communications
iOS	iPhone operating system
IoT	Internet of Things
IP	Internet protocol
IT	Information technology
LPWA	Low power wide area
LTE	Long Term Evolution
MNO	Mobile network operator
MTC	Machine type-communication
MVNO	Mobile virtual network operator
OTT	Over the top
P2P	Peer to peer
PLC	Power line communications
SCADA	Supervisory control and data acquisition
SDK	Software development kit
SIM	Subscriber identity module
UMTS	Universal Mobile Telecommunications System
V2P	Vehicle to Pedestrian communications
V2V	Vehicle to Vehicle communications
WRC	World Radio Conference

XONA PARTNERS

Riding the Advanced Cloud Deployment Roadmap

Creationline, Inc. Team

Dr. Riad Hartani, Rolf Lumpe (Xona Partners)

August 15th, 2014

Table of Contents

1	SYNOPSIS	3
2	RATIONALE FOR A CLOUD INFRASTRUCTURE ROLLOUT REVISIT	4
2.1	Cloud Migration Considerations	4
2.2	Evolution to Advanced Cloud Infrastructure – Challenges	5
3	CLOUD INFRASTRUCTURE – THE ROAD AHEAD	6
3.1	Understanding the design of the IaaS / PaaS component, as a leverage into migration trade-offs analysis.	7
3.2	Migration	12
3.3	Monitoring, Diagnostic and Action models	13
3.4	Optimization	13
4	CONCLUSIONS & CALL FOR PARTNERSHIP	15
5	ACRONYMS	16
6	THE CREATIONLINE TEAM	17
7	THE XONA TEAM	18

1 Synopsis

The last few years have seen a select set of large scale Internet players position their Information Technology (IT) infrastructure as their core asset over which their services would ride, and as such have, for their in-house needs developed a lot of what would form the basis of the cloud as we know it today, with most of it coming out of their labs in the form of Open Source software and solutions – This including the various aspects related to IaaS, PaaS, Big Data and Data Sciences. Most recently and something that we will very likely witness over the next few years, a progressive move towards the use of cloud solutions will be the norm for a variety of corporations and service providers. This will be gradual, this will be on a need basis and this will be function of technology maturity and foreseen returns on investments. More importantly, this will require the emergence of advanced cloud centric product and services teams that could assess such migrations, develop them, deploy them and support them.

This is exactly where, Creationline, Inc. (“Creationline”), in collaboration with Xona Partners (“Xona”) have set sights in terms of putting together a cloud specific transformational information technology practice to address these upcoming challenges.

Starting from our Tokyo, Japan and San Francisco, California’s Silicon Valley head offices, we set sail for a journey around what we see in terms of Cloud Infrastructure evolution challenges, and highlight our evolving contributions aiming at overcoming them. This short positioning paper, is presented as a baseline for follow up detailed discussions related to the various topics under consideration, with the goal of designing, customizing and optimizing our solutions to lead data centric organizations’ needs, leveraging the broad and complementary expertise of our team.

Specifically, the paper presents a comprehensive methodology, which includes the assessment of various models for cloud migration, the design and implementation related to building IaaS/PaaS/SaaS and porting applications to these environments, as well as the operational procedures required for a successful completion of cloud design projects. Along with this, an innovative cloud monitoring and optimization solution is introduced, with the aim of dynamically adapting cloud resources, based on the processing performance requirements.

2 Rationale for a Cloud Infrastructure Rollout Revisit

The fundamental premise of the question we are addressing is fairly simple:

As a business that is built around gathering large-scale data sets, mining and learning through such data, and optimizing communication between those producing it and those using it, where shall I go from here? This would touch upon the evolution of the compute, storage and network platforms over which data sets sit and business services processed.

Most information technology players, are, or will soon be considering enhancing their IT and data management infrastructure to help build data solutions which will optimize their ability in identifying, capturing, and managing data to provide actionable, trusted insights that improve strategic and operational decision making, resulting in incremental revenues and a better customer experience.

2.1 Cloud Migration Considerations

The current challenges of the existing platforms mostly affect the operations teams' ability to provide reliable SLAs for the reporting jobs that are critical for the business, the data platforms required scale, and perform effectively as the business grows. As such, the desired goal is to create a solid foundation architecture that is able to provide these optimal functional capabilities, and a platform to overlay additional applications such as intelligent business intelligence and Data Science as a service capability.

For most players, an insider-only approach, leveraging internal resources, would only go so far in terms of architecting and designing this new cloud based architecture, given the breadth in terms areas of expertise required, and more importantly, the need for a step-back and outside the box design way of doing things. As such, our team aims to be the strategic partner bringing in such expertise, and build open models tested and validated over a large set of data platform development models.

It is worth observing that:

- The expertise in deploying large scale cloud infrastructure still sits in the hands of the large cloud players themselves, and being able to predictably design it and build it, would most likely require the expertise of teams who have done it for their own IT within these large cloud players
- The Open Source software (such as OpenStack, CloudStack, Cloud Foundry, etc.) is still in constant mutation and likely to evolve fairly rapidly over the next few years, requiring specialized teams to bring in to market, deploy it, evolve it and manage it.
- The most crucial component of cloud migration is to figure out the right ROI model, based on which applications are migrated, how they are migrated and how they are used post migration. This in turn makes it primordial to figure out the right cloud model (public, private, hybrid) as well as the right framework to monitor such cloud deployments when done.

Given the above considerations, one can see that we are in the early days of large scale cloud migration, and the specialized expertise to do so is what is needed the most at this point in time. As such, we present our views on what methodology shall be implemented and what is likely to become the cornerstone of any related future cloud rollout roadmap.

2.2 Evolution to Advanced Cloud Infrastructure – Challenges

As of today, the existing information infrastructure and analytics processes suffer from challenges we have observed and worked on in the most typical large scale data and IT projects, some of which are listed below:

- The requirement to first progress the overall virtualization of computing and storage, and increasingly network resources. The increasing tie up between the virtualization and IaaS/PaaS environments renders the virtualization strategies very much dependent on a more forward looking broader cloud migration strategy.
- Understanding the variety of software applications and service running within the IT environment and analyzing them individually, then as an aggregate as far as ways of evolving them towards a cloud model
- Understanding the still evolving toolsets for cloud services monitoring and diagnostics of distributed software applications and compute processes
- Requiring a coordination between what would run in-house as a private cloud, or externally hosted private cloud off premise, with what runs over chosen public clouds
- Understanding how internal business processes shall evolve to accommodate such cloud migrations, and in fact, having that as a constraint that would force specific directions in such migration.

In the upcoming sections, we position our methodology, on how these cloud architectures should be evolving, short, mid and long terms. The analysis is presented in a generic fashion, but builds on top of a very selective and specific set of case studies that we have worked on in the real world, developed and completed. In other words, what is described is a pragmatic successfully completed case study, albeit made generic, to show broader applicability.

3 Cloud Infrastructure – The Road Ahead

Our approach to tackling evolutions towards cloud architectures is based on our understanding of the underlying business models, the existing IT and data architectures and deployment patterns and the assumptions we set, as far how data information models are structured and the underlying performance and reliability requirements. Additionally, our methodology and solutions aim at providing a capability maturity path for increasing capabilities of the system with minimal disruption to current operations, forming the basis for an evolutionary migration. As such, the architectural models we built our platforms upon, are designed to address some of the persistent problems in current infrastructure and future needs of the organization for data processing.

The figure below shows a standard component of a data center (DC) cloud infrastructure, including the compute, storage and network components.

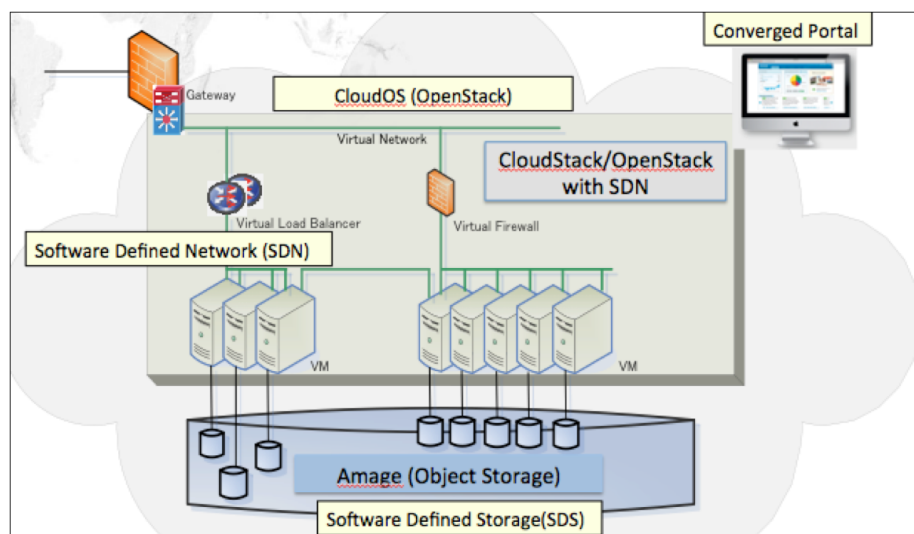


Figure 1: Standard components of DC cloud infrastructure

The figure below shows a complementary perspective, related to the various layers of the cloud implementation within an IaaS provider, and the various components required achieving it.

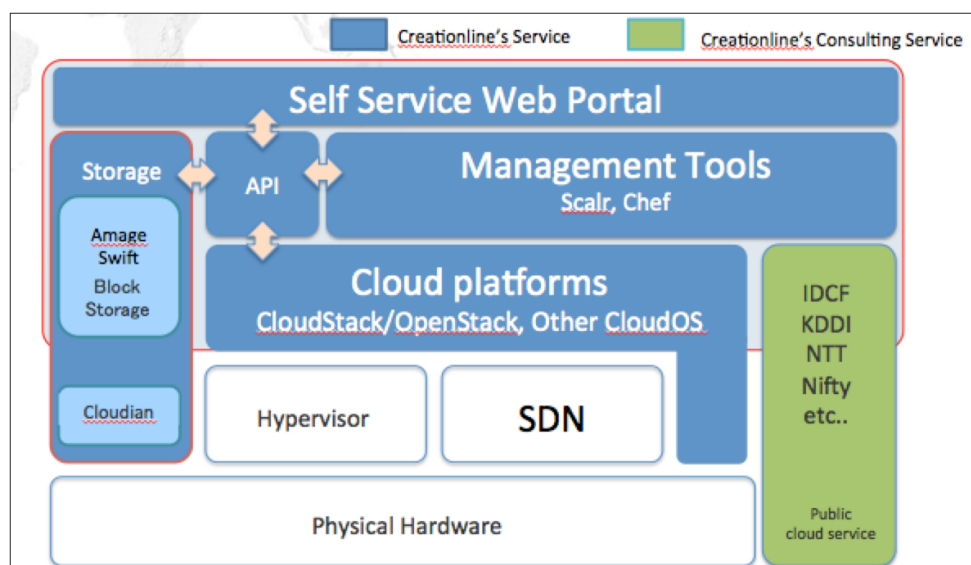


Figure 2: Cloud implementation for IaaS provider

Our methodology is based on four distinct steps. The first step is to have the target cloud infrastructure architected and implemented to allow any migration towards such target. Such design, which would form the basis of the IaaS and PaaS components, would be done with the right public or private cloud provider. The second step would approach the problem from the customer side of the cloud infrastructure, in other words, the corporation owning the IT and services that would want to run them over a cloud model. This would lead to the assessment of what would migrate, how and when based on the understanding of the IaaS and PaaS designs done prior. The third step would focus on supporting the migration and post migration, which would include performance and fault analysis and diagnostic. The last step would focus on overall optimization via appropriate orchestration and automation, and leveraging arbitrage models to select what to migrate on which cloud over time based on ROI dynamics.

The figure below shows the components of the methodology. It includes a non exhaustive list of technology partners with whom we are working to put such methodology into practice. Such eco-system is rapidly evolving and would need to cover the various options required to satisfy the diverse needs of cloud migration initiatives.

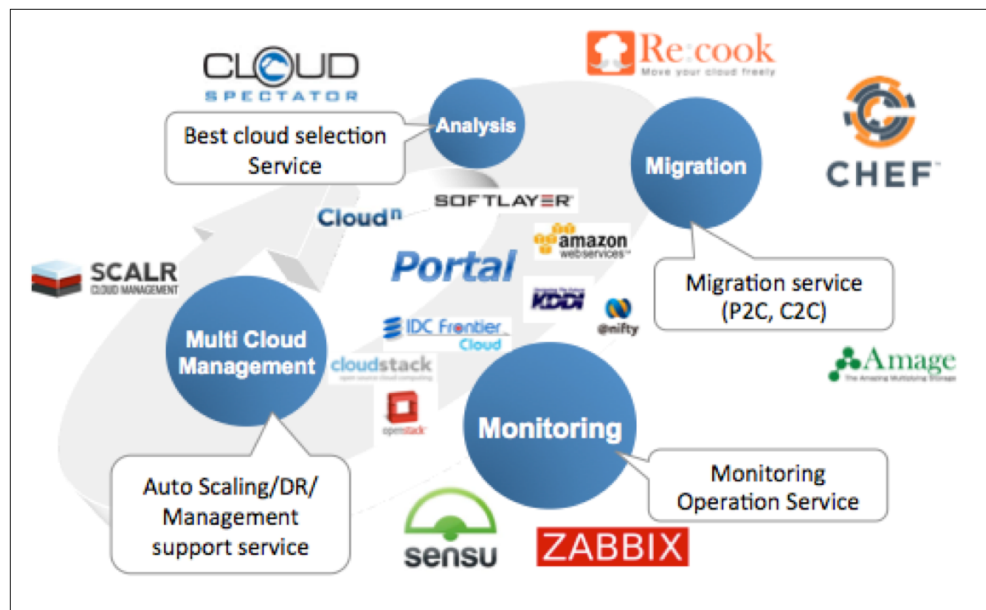


Figure 3: Components & methodology

We briefly describe these 4 steps and highlight the key aspects to look into while achieving them.

3.1 Understanding the design of the IaaS / PaaS component, as a leverage into migration trade-offs analysis.

Three pre-requisites are required to implement an adequate cloud migration strategy. First is a comprehensive understanding of the target virtualization, IaaS / PaaS environments, second is a tight working relationship with the cloud eco-system players into this environment and third an understanding of the IT and application software environment that would be candidate for migration and specifically its big data management component, as this would, in most cases, form the basis of the business ROI when migrating application and data management to the cloud.

These three aspects are briefly described.

Understanding The Virtualization, IaaS and PaaS designs and engineering considerations

As applications migrate to the cloud, the first thing to do is to figure out what to migrate, to what and how. Our thesis is that the best way to provide the best answer is to first have a detailed understanding of the design of IaaS and PaaS platforms and overall virtualized environments in the DC. Such knowledge can only be achieved through having led the design of these virtualization, IaaS and PaaS deployments.

Our experience having led the rollout of large IaaS solutions, with both CloudStack and OpenStack is what we build on to provide insight into what target IaaS and PaaS models would be optimal for various migrations, as well as laying out the right engineering and implementation models.

The figure below describe a high-level implementation example, where a CloudStack based IaaS and PaaS have been designed from the ground up to host enterprise applications in a telecom provider public cloud infrastructure.




Figure 4: High-level CloudStack implementation (host enterprise applications)

In a similar way, below describes a high level OpenStack IaaS/PaaS implementation over which a high availability storage solution, complemented by a hybrid cloud disaster recovery solution as a service has been implemented in a Tier cloud services provider.




Figure 5: High-level CloudStack implementation (high availability storage)

Understanding the evolution of the partner eco-system

Understanding the various options and trade-offs in building the cloud infrastructure would require an eco-system of partners, where insights and rollout experiences are shared and contrasted. Below is a description of the partner solutions hierarchy, as well as illustrative parties within such eco-system that we have been working with. It is worth noting that most of these partner solutions' focus is in deploying, supporting and commercializing open source solutions. This list is fast growing with most partners being less than few years old, which is a testament of the novelty of the whole industry.

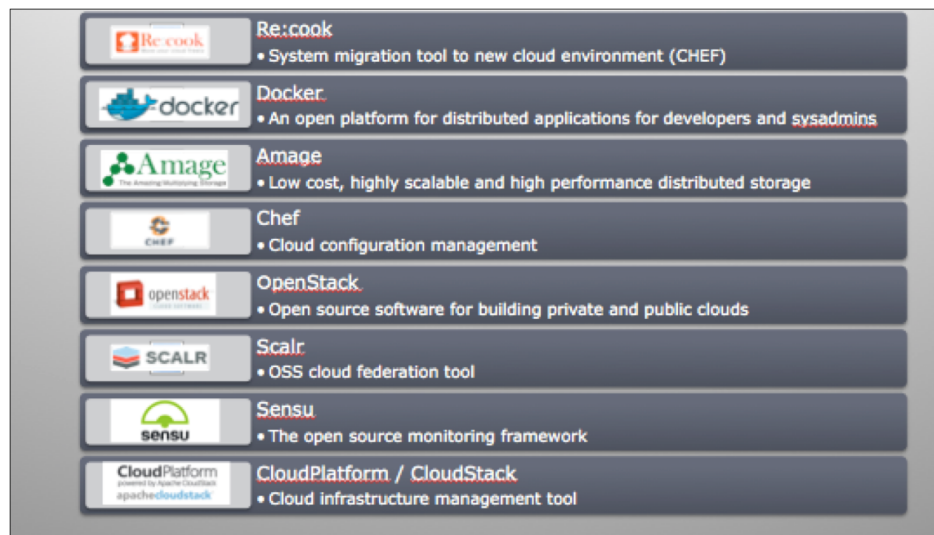


Figure 6: Partner Solutions

Understanding the Big Data Sciences Angle

Big Data Sciences is a combination of technology, business and mathematics that increasingly

impacts every facet of daily life. The combination of traditional disciplines of data extraction, data intelligence, data analytics, data modeling, data warehousing, and reporting along with statistics and predictive analytics can be referred to as Data Sciences as illustrated in the diagram below.




Figure 7: Big Data science angle

Overall, the Data Science requirements would direct link to the data management (based on open source or commercial frameworks such as Hadoop or the various SQL variations) and analytics derived from such data sets, which in sits on top of the cloud infrastructure. This in turn sets a lot of the requirements that one would have tackle in designing such infrastructure, based on the recursive logic of first understanding the end user application, the underlying data management with impact on the cloud implementation. A natural way of visualizing the various components of the Data Sciences hierarchy is shown below, taking it from data extraction at the bottom all the way to applications to specific verticals at the top.

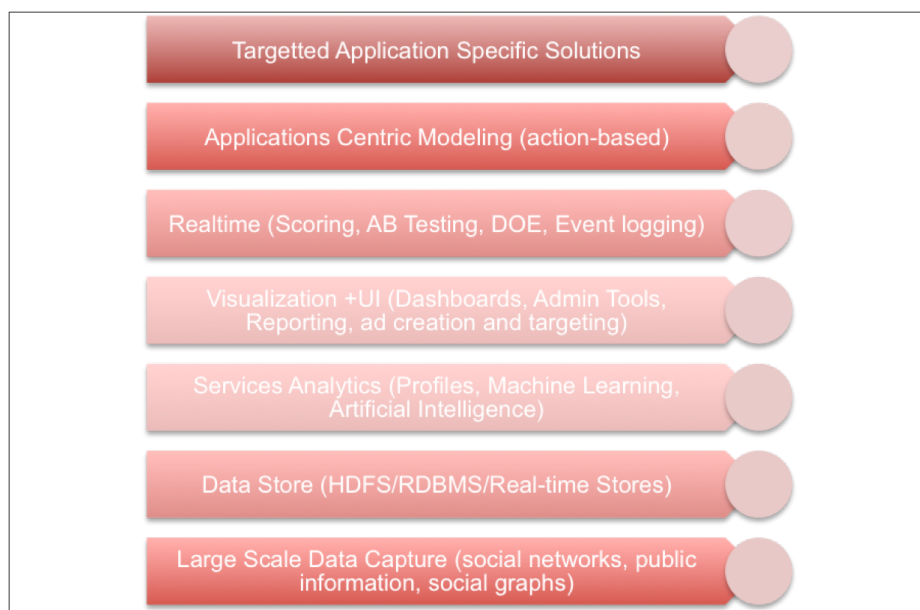


Figure 8: Big Data science angle

Based on the framework, which we have developed and put into practice on real case studies, specific credit transformational models will be highlighted. Below is an illustration of such cloud migration in the online advertising space where data is aggregated, processed and exposed by real time bidders for online ads (as Demand Side Platforms, Sell Side Platforms or Data Management Platforms). This is implemented over multiple steps, including the porting of data management tools into a Hadoop based management platform, which would sit within a private/public cloud environment.

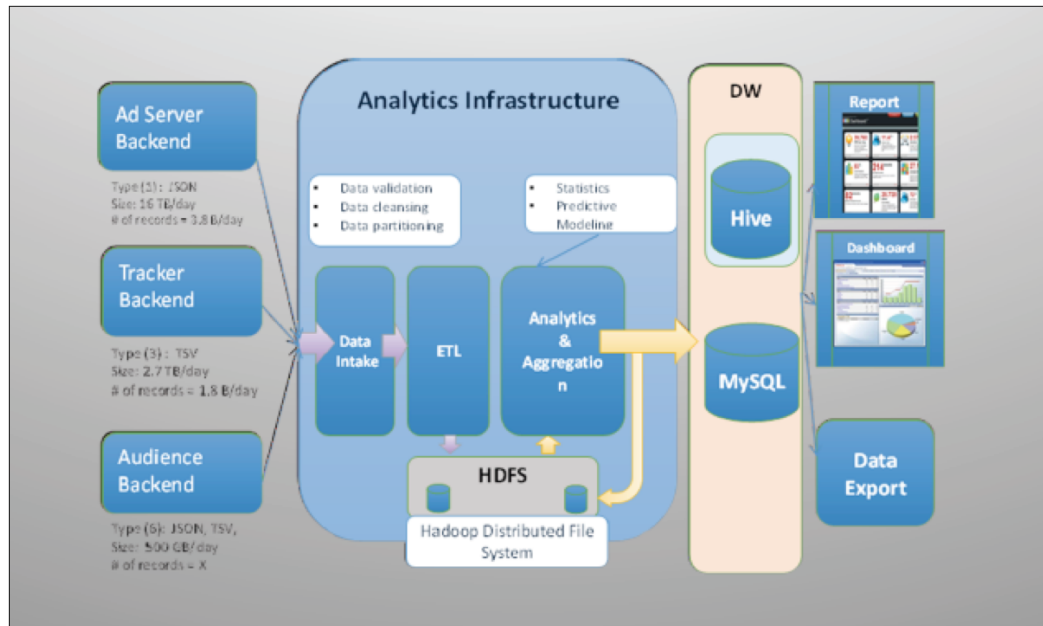


Figure 9: Analytics Infrastructure

It is this understanding of (a) the IaaS / PaaS environment, (b) the cloud eco-system players into this environment and how they are evolving, and (c) the big data and data science angle that would form the basis of the business ROI when migrating application and data management to the cloud that would form the cornerstone of the overall analysis.

3.2

Migration

Taking into account the specifics of the various private and public cloud models that one could migrate to, the next step is to understand the migration of the various applications and analyzing the underlying trade-offs, as described below:

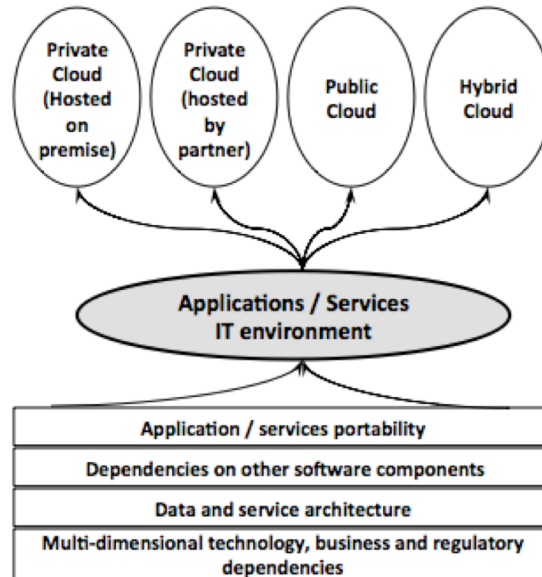


Figure 10: Migration methodology

This methodology shall result into a recommendation model in terms of migrating various applications based on their intrinsic requirements, as illustrated below:

Evaluation Criteria	Private Cloud (Premise Hosted)	Private Cloud (Partner Hosted)	Public Cloud	Hybrid Cloud
Agility / Migration	Medium	Medium	Medium	Low
CAPEX / OPEX	Medium	High	Low	Medium
Application Portability	Low	Low	High	Low
Security / Privacy	Low	Medium	High	Medium
Operational Complexity	Low	High	High	Medium
Other Criteria (function of context)	Ranking	Ranking	Ranking	Ranking

Table 1: Migration model

3.3 Monitoring, Diagnostic and Action models

Once migration is executed, the next step is to monitor its evolution and analyze the live performance, reliability and quality of service requirements. This would include the deployment of monitoring tools, interfacing with the resource management and orchestrator tools and having access to preventive and corrective action models, primarily via the automated provisioning tools, as illustrated below.

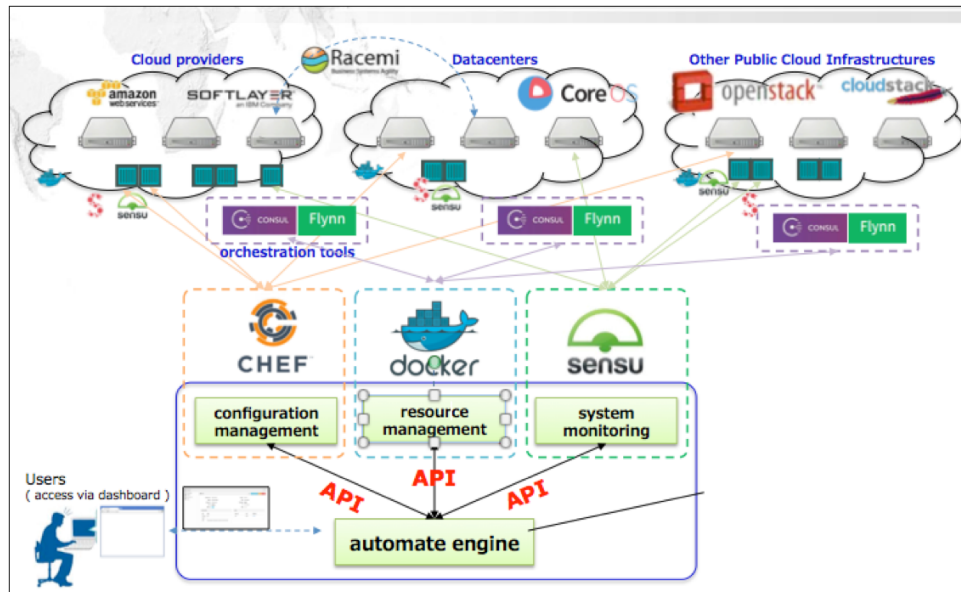


Figure 11: Cloud intelligent monitoring engine

The monitoring functionality is an essential feature for any system relying on various degrees of automation. The various system components work together utilizing the monitoring and automation APIs. Given that the various cloud infrastructure platforms and tools support the RESTful API, The proposed includes a comprehensive monitoring system, which dynamically interfaces with the various other cloud management tools. This intelligent monitoring and automation tool, which has been developed, based on the most stringent reliability and performance requirements of some of the most advanced cloud environments is described in the following section.

3.4 Optimization

As the application are ported in the cloud, and the multi-layer monitoring, including the IaaS and PaaS environment, as well as the application software and services environment, the last step is to put in place a dynamic optimization model that would analyze next actions to implement. This would include actions for optimizing reliability and predictability, as well as overall cost dynamics leveraging the choices offered by the various players in the eco-system, including data centers and cloud providers, systems integrators, support partners and the remaining players. This is mostly done via a cost arbitrage function that would dynamically recommend ongoing application migration options to select cloud environments, as described in the figure below:

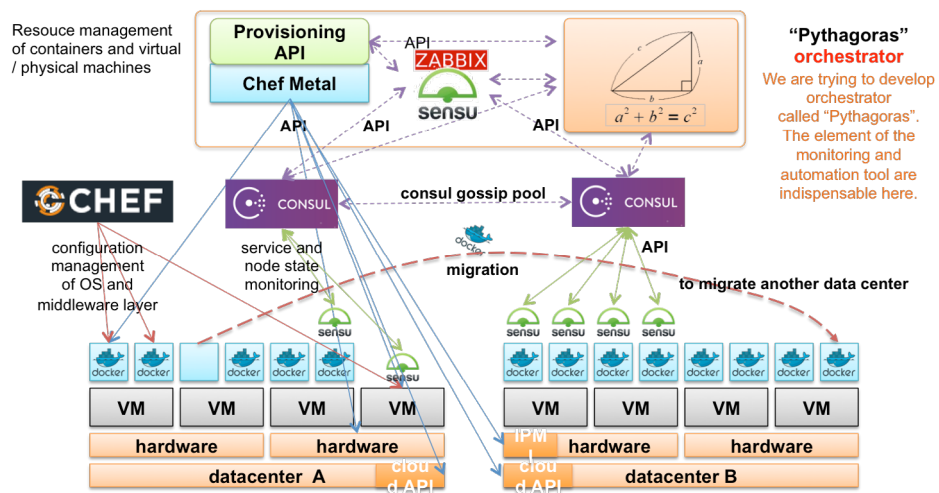


Figure 12: Dynamic optimization model

The above illustration describes an architecture where the cloud compute, network and storage resources are managed by the intelligent orchestrator, names “Pythagoras”. This not only includes the datacenter underlying infrastructure but also the hardware resources, virtual machines’ resources and “Docker” container resources. The orchestrator integrates “Chef Metal” and “Consul” for resource management and execution, based on the active monitoring and the respective data acquired by “Sense agent and APIs”. The “Pythagoras” orchestrator enables a highly reliable and a fully optimized automated cloud operation.

4 Conclusions & Call for Partnership

Following some successful validation over the last few years, where some of the largest cloud and big data transformation projects have been conducted, involving designs with some of the most aggressive scaling, reliability and manageability deployment requirements, we are now in the process of taking in-house development and deployment methodologies to the larger market, and would welcome discussing specific requirements with key IT and cloud solutions architects having for a mission to lead their IT transformation architectures, as well as with managed services players wanting to build on their existing IT and big data capabilities and augment it with specific cloud based data management platforms.

Specifically, the analysis models developed for cloud migration assessment, the software tools and processes in use for the IaaS/PaaS and SaaS design and operations, as well as the intelligent monitoring, automation and orchestration tools for an optimal cloud operation would provide us with an optimal starting point in analyzing the specifics of the cloud architecture, and ensuring its successful deployment.

We believe that the next generation cloud and big data platform architectures will be evolving in the direction we have been highlighting, and hence, encourage various players to speed up such evolution, for the common interest of the various eco-system players.

5 Acronyms

DC	Data Center
HSFS	High Sierra File System
HDFX	Hadoop Distributed File System
IaaS	Infrastructure as a Service
IT	Information Technology
PaaS	Platform as a Service
RDBMS	Rational Data Base Management System
ROI	Return On Investment
SaaS	Software as a Service
SLA	Service Level Agreement
UI	User Interface

6 The Creationline Team

Creationline was founded in 2006 and developed over the past 8 years as Japan's most advanced cloud enabling and professional service company. Initially providing consultation services to major Japanese telecom carriers and IDC firms building cloud service infrastructure based on open architecture Cloud such as CloudStack and OpenStack, as well as big data architectures such as Hadoop, and large scale PaaS infrastructures such as Cloud Foundry. Services include Proof Of Concept (POC), design, implementation and support. Nowadays Creationline service set includes migration (P2C, C2C), monitoring & operations, multi-cloud management and cloud building services.

Sample relevant globally recognized projects completion includes Softbank Japan Cloud Services Design, KDDI public and private hosted cloud services design, NTT Cloud Services architecture as well as various cloud & Big data engagement with select data center providers.

XOONA PARTNERS

Strategic Advisory The Case for a Disruptive Model

August 2014

Synopsis

Xona Partners (“Xona”) team members have been heavily involved in developing various advanced technology and business innovation models, and observing what works, what doesn’t, and why. After analyzing the latest shifts in the technology eco-system and the competitive positioning of lead technology players across key markets around the world, we have reached a startling conclusion: traditional consulting and advisory models are no longer optimal for the needs of leading edge technology businesses.

We are pioneering an execution-driven approach based on a shared risk and shared return model and focused on accelerating innovation disruption to create entirely new value chains for our clients. This model is based on a technology incubation approach, followed by a progressive spin-in into the client’s business, creating as such new revenue streams in adjacent businesses.

In this paper we present the underlying rationale, highlight its fundamental components, and illustrate specific case studies conducted in the US, Japan, Korea, Hong Kong and Singapore of how this execution centric technology advisory model has been implemented in partnership with some leading technology companies and private equity investment firms.

Rationale

Collectively as a team, we have spent the last 2 decades fully immersed in various innovation ecosystems around the world, and we have approached the various aspects of technology innovation from different angles. This includes jumpstarting venture capital-funded startups and taking them to acquisition or public markets, building new businesses out of corporate and academic R&D work, assessing and executing merger and acquisition (M&A) for technology businesses, working with boards of directors on their business strategies, advising investment and private equity funds on their technology investments and management of their portfolio companies as well as directly leading angel and venture capital investments.

Our involvement has not only been focused in the Silicon Valley area, where most us built their technology startup roots, and where a lot of the technology disruptions in information technology (IT), Internet business models, cloud and data sciences are still happening. We have also participated in direct engagements within various innovation ecosystems in Japan, Korea, China, India and Europe, among others. These innovations have also been taken to markets in both the developed and the emerging world by Xona, are now running commercially with a validation of both the technology and the business models.

Over the last few years we have honed our methodology for assisting various technology businesses, technology investors, government organizations and policy makers in developing & executing new business models. This has primarily been a response to the observations made above: the inadequacy of the traditional advisory model. Why do we say this?

- The pace at which technology and information technology is progressing has shifted: this is leading to a brutally fast disruption of existing business models, accelerated convergence, and shifting revenue and margin dynamics between competing businesses. As an illustration, an observation of the market cap of various large technology businesses shows how disruptions happening over a timeframe of a few years can lead to a total reconstruction of the leading pool of players, from mainframe businesses in the 80s to networking vendors in the 90s

to cloud and Internet players in the 2000s, and progressively into media and e-commerce players moving forward. This fast evolutionary pace has put pressure on boards and CxO's to increase risk tolerance and accelerate the speed of decision-making in terms of what to build and which adjacent market to target for expansion.

- The convergence of various industries driven mainly by the way information is exposed, exchanged and consumed, has led to the need to assemble a very diverse expertise to tackle the problems associated with new business creation. The need for technology-centric commercial expertise which can link into new business growth and which is readily available at short notice is growing and represents a clear value differentiator for decision makers.

- Successful integration of disruptive business models is hard. Apart from a very small number of technology players that have mastered the art of startup acquisition or larger business integration, few organizations have been successful in delivering disruptive value, forcing them to develop most of their new products in house with direct impact on the likelihood of success, bottom line and time to market. The Internet business model, where the "winner takes all," is likely to exacerbate this trend.

- The traditional technology and business strategy advisory model has remained as it was in the 1980s, 1990s and 2000s: focused on high-level strategy without direct coupling with operational execution, and focused on short-term technology and business trends without a deep dive into the implementation of technology to understand its impact on the creation of new businesses in adjacent markets. While this traditional model is still effective in some situations, it has not been successful in adapting to the changing information technology environment.

Given all the above, we believe that the speed of change in technology and the variety of options available for boards and management teams in expanding into adjacent businesses, requires an immediate access to deep technology expertise, combined with both operational know-how and strategic understanding of business implications are primordial. We have developed a model that we believe accommodates these requirements as it forms the basis of a new technology and business advisory practice.

The Advisory Model Explained – A Technology Spin-in & Incubation Centric Approach

Our advisory model has been developed and validated over the last 2 years in close partnership with lead technology groups aiming at expanding into adjacent businesses, as well as private equity groups focused on creating new value out of their portfolio technology companies through new business creation. It is based on the following guiding principles:

- Strategic Technology Incubation

By working with key decision makers (typically, the business/strategy and engineering leads) within a technology corporation, our team will create a collaborative plan to develop a specific solution focused on target adjacent markets, leveraging our specialized hands-on design expertise and business insight for new market insertion. Such know-how will progressively be transferred to in-house engineering and sales teams, via a tailor made enablement process. This would form

the basis of the incubation process.

- Integration via a “Spin-In” Model

Our team will have as a goal, to first incubate and progressively develop a solution within a 12 to 18 months period, based on pre-agreed and designated milestones with the business stakeholders. This is primarily used to de-risk decisions for decision makers prior to committing broader company resources, and progressively build the required expertise to take over execution. The incubated solution will be integrated into the technology corporation’s mainstream process via a spin-in model.

- Progressive Integration of Synergetic Growth Models

The spin-in model is based on a milestone-driven approach. During this process, various strategic alternatives may be selected to penetrate these adjacent markets: This includes recommendations of possible technology acquisitions, in a buy vs. build model, a fast tracked strategic investment to speed up an existing development process, or an asset carve-out strategy leading to a more optimal business strategy. These various alternatives will be analyzed, contrasted and if relevant, recommended by our team as an alternative to a technology incubation and spin-in approach.

- Adjacent Business Growth Via Shared Risk and Shared Return

The fast pace of information technology innovation, leads to a large number of business alternatives to choose from with the goal of entering new adjacent markets. These choices come with a significant business risk as well as a high execution cost. As such, a shared risk shared return model is the most viable approach for decision makers. This approach mimics the technology startup model, and creates an incentivized environment. Our team would be sharing the innovation and execution risk with the business stakeholders as a way to de-risk their decision to get into adjacent businesses.

Our technology advisory practice provides clients with an exclusive capability: the ability to deliver innovation and disruption in a risk-mitigated and value-optimized form. This model, aligned with the few leading technology companies who have successfully achieved this with in-house teams, enables organizations to maximize incubation and spin-in approaches. Our aim is to make such a model available to a larger set of technology players based on our team’s diverse experience and expertise, risk and return sharing DNA and focus on operational execution.

This model is valuable to a set of stakeholders, and primarily to:

- Advanced Technology Organizations

Such corporations represent the primary beneficiary as they are, and will increasingly be, likely to be expanding into adjacent markets and building corresponding large businesses. As such, they would be the primary partners in the proposed risk/return-sharing model.

- Technology Investors (Private Equity, Venture Capital)

As shareholders in the various businesses they invest in, the benefits from this model are clear. Assisting portfolio businesses in either developing new businesses in-house via the spin-in

model, or moving into tackling growth via an M&A model, or in other cases, executing technical & commercial due diligence, converging on asset carve-outs and undertaking strategic restructuring for better long-term synergetic growth.

- Policy Makers (Development of Innovation Eco-Systems)

Along with technology businesses and investors in technology policy makers responsible for bringing innovation into their own ecosystems would also be beneficiaries of our innovation advisory model. Specifically, this would be via working with regulatory arms and investment groups to better position new technologies for the specific needs of the ecosystem, and therefore enabling the emergence of the appropriate climate for innovation with direct implications on the development of new technology markets and businesses.

To better illustrate this model, some use cases will be briefly discussed with highlights of the value proposition to the stakeholders.

Adjacent Growth Businesses Incubation - Case Studies

We have experimented with several case studies. Most of these engagements are described in select whitepapers (1). Overall, all of these engagements have in common the following characteristics:

- The team brought immediate technology and hands-on expertise that wasn't readily available to the client.
- The team worked with the business stakeholders, either management or investors, to analyze various technology trends and associated business impacts and zoom-in on a select set of new adjacent markets to address. These markets would be of a Greenfield nature for the client, and would take 12 to 18 months to bring to market commercially.
- Addressing these markets would still consider alternative approaches to the in-house spin-in incubation, such as M&As or business assets carve-outs.
- The team formed a task force as an extension to the client's team, to lead technology, business development, customer engagement, sales enablement and solution validation with lead customers
- The team defined milestones jointly with the client, and compensation was on a shared risk/return basis, in a very similar way to the technology venture capital startup model.

Select use cases that we have conducted over the last 24 months, in North America, Japan, Korea, Singapore and Hong Kong are briefly presented below to illustrate the work done in various technology areas. The respective references point to detailed information describing how these projects have been executed with various technology clients.

- a) Data center hosting providers have a clear need to evolve towards hosting new technology eco-systems. As such, we put together an architecture, design and execution model for the development of a new revenue generating business based on hosting the fast growing ecosystem players in real-time bidding for online advertising (2).
- b) As various businesses explore ways of leveraging the availability of vast amounts of data

as well as big data frameworks to manage it, we developed data science centric solutions to create new revenue streams from the analysis of these data for various adjacent markets to specific industry verticals, including mobile payments, mobile analytics and vehicular technologies (3).

c) The evolution of mobile health, in conjunction with the emergence of more robust health centric wearable devices, opens up the opportunity for mobile and virtual operators. We developed and implemented a new health vertical for mobile players (4)

d) The emergence of cloud delivery models opens up interesting disruptions for various industry players, and new entry points into adjacent businesses. We architected a solution leveraging open source models as an entry point into the hybrid cloud service business, complemented with tailored IT and big data transformation services (5).

e) Telecom operators are under increasing pressure to optimize capex and opex models. Sharing infrastructure is key to achieve savings. We have developed and implemented new solutions for active infrastructure sharing and developed a baseline for partnership with OTTs as MVNOs (6).

f) Financial technologies are under pressure to leverage new IT and data science models to optimize their bottom lines. In this context, we have developed new technology and business solutions around data sciences for financial industry players with a focus on integrating intelligent automation into wealth management in this context (7).

g) As various eco-systems around the world aim at leveraging the advantages of digital economies, a race towards the creation of technology innovation hubs, blending academic, private and public capital funding has been occurring over the last decade. Building on our methodology of incubating new businesses, in partnership with the various players in the eco-system, we have partnered with different corporate groups in a way to best synergize with the already successful Silicon Valley innovation eco-system (8).

Conclusion

We have highlighted the rationale for a drastic change in the way strategic technology and business advisory is being conducted, in a world where fast changing Internet & Information Technologies is at the center of innovation, leading to drastic business disruptions. The fundamental reasons for such disruptions are highlighted, forming the basis of a new model, which we have developed, tested and validated.

This model is primarily built on our broad Silicon Valley technology startup culture, with a shared risk and return philosophy. It leverages the disruptions seen in the information technology world, as far as rate-of-change, industry transformation, and business models mutation, and is based on a technology incubation model, complemented by an operational focus and a shared risk return execution philosophy. Through the recent successes in deployment with various technology businesses and investors in technology, we believe this represents a near-optimal approach for technology and business advisory in the near future.

References

(1) Select Xona business incubation and spin-ins white papers

<http://www.xonapartners.com/media/whitepapers/>

(2) Development of a Online Advertising Eco-system for Data Center hosting operators

<http://www.xonapartners.com/wp-content/uploads/2014/01/PrivateCloudforOnlineAdsRTB.pdf>

(3) Development of data science centric adjacent businesses for select industry verticals

<http://www.xonapartners.com/wp-content/uploads/2014/01/DataScienceAPracticalPerspective.pdf>

(4) Development of an e-health eco-system for mobile operators

<http://www.xonapartners.com/wp-content/uploads/2014/01/DigitalHealth.pdf>

(5) Development of an IT transformation Service business based on cloud and big data technologies

<http://www.xonapartners.com/wp-content/uploads/2014/01/DataManagementInfrastructureTransformation%E2%80%933PositioningandProposal1.pdf>

(6) Development of new models of active infrastructure sharing in wireless networks

<http://www.xonapartners.com/wp-content/uploads/2014/01/SpectrumAndNetworkSharingModels%E2%80%933TrendsBusinessImpacts1.pdf>

(7) Development of data driven intelligent advisory software in financial technologies

<http://www.xonapartners.com/wp-content/uploads/2013/05/DataSciencesinWealthManagement.pdf>

(8) Synergizing with the Silicon Valley innovation eco-system

<http://www.xonapartners.com/wp-content/uploads/2014/01/SiliconValley.pdf>

XONNA PARTNERS

Innovations in RF Distribution Networks: Evolution of Distributed Antenna Systems

By Frank Royal

July 28, 2014

Table of Contents

Overview	3
Market & Technology Drivers	3
The Options	5
Evolution of DAS Systems	9
A Perspective	9
The Present	10
Evolving Trends	12
Enabling New Business Models & Applications	15
Conclusions	16
Acronyms	17

Overview

The market for in-venue communication systems continues to expand steadily with the promise of accelerating growth in the future. Distributed antenna systems (DAS) have been the primary type of system deployed in venues, but alternatives are available on the market. While this growth is a direct response to ever increasing demand for mobile data services, there are a number of trends that combine to shape and influence the development of this market in the short and medium term. In this report, we seek to identify the trends that shape the market for in-venue communications with particular focus on the evolution of DAS and its outlook over the next 2-5 years. We also argue that it is the business model and applications enabled by DAS and competing technologies, as well as operators' attitude towards such business model that would shape the outcome of the competitive landscape. Seen from this perspective, DAS has had the advantage of allowing operators to share a common infrastructure. The evolution of DAS provides for new applications and opportunities that are outlined below.

Market & Technology Drivers

Demand for mobile data service concentrates indoors and in venues¹ where as much as 85% mobile traffic is generated. Subscriber behavior enabled by the proliferation of smartphones and other types of mobile computing devices, such as tablets, coupled with social networking applications are especially bandwidth consuming. For perspective, data traffic first exceeded voice traffic on mobile networks at the end of 2009 when traffic was 100 petabytes per month. At the end of 2013, traffic ran at 2,000 petabytes per month and is expected to surpass 15,000 petabytes per month in 2018. Wireless network performance cannot help but be adversely impacted by such high localization of indoor traffic because of factors such as propagation losses into structures as well as high oversubscription to limited capacity resources. Placing wireless transceivers at the venue becomes mandatory as mobile network operators (MNOs) look to improve service performance in the venue as well as to free adjacent cell sites covering the venue from a singularly demanding traffic hotspot. This has been, and continues to be, the primary motivator for in-venue solutions – a market valued at \$10 billion in 2018². However, trends in mobile communications targeted to improve broadband data services as well as divert traffic away from loaded macro cell would increase the demand for in-venue solutions.

To validate our position, consider the following:

1. It is more difficult to penetrate buildings with broadband wireless coverage than narrowband coverage. Wide channels have reduced coverage footprints and lead to a shorter range of service in comparison with narrow channels (Figure 1). This becomes more acute in technology like LTE where the channel bandwidth reaches 20 MHz, or 4 times that of 3G and 100 times that of GSM. While LTE does include other techniques that reduce some of the lost system gain due to channel bandwidth such as convolutional turbo codes, multiple antennas, and hybrid ARQ, these techniques do not combine to improve capacity where the communication link is weak.

¹ In this paper, the term venue refers to a high concentration of subscribers indoors or outdoors in facilities such as stadiums, airports, train stations, campuses, large commercial buildings, hospitals, etc.

² Mobile Experts, "Mobile Experts Identifies \$100B In-Building Wireless Infrastructure Opportunity," April, 2014.

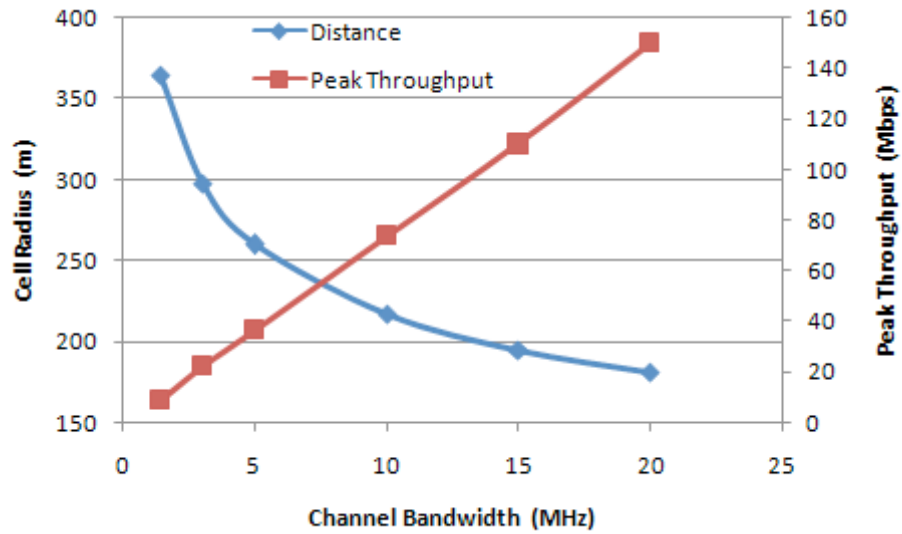


Figure 1 Distance and peak throughput performance for 2x2 MIMO LTE micro cell in urban clutter.

2. The challenge of serving venues is increasing in magnitude as regulators release spectrum in higher frequency bands for mobile service such as 2300, 2600 MHz, and soon 3500 MHz as in Japan by the end of this year. Propagation and wall penetration losses increase with frequency, resulting in consecutively smaller coverage footprint for higher frequency bands.



Figure 2 Coverage distance for different spectrum bands.

3. High throughput requires high signal quality. Efficient modulation such as 64QAM (6 b/s/Hz) and MIMO spatial multiplexing necessitate high signal to noise and interference ratio, for example, exceeding 18 dB. The ability to achieve the high signal quality and level required to engage these features degrades as the signal attenuates upon entering the venue.

In summary, the emergence of LTE coupled with the drive to supply ever higher capacity to concentrations of subscribers in venues is set to accelerate the in-venue communication market. Operators view such venues as strategic service locations which they cannot easily surrender service within to a competitor. Complementary to this, serving a traffic hotspot venue is a means to offload key cell sites of traffic and allow them to operate for their intended service target.

There are several options for mobile network operators to provide service in venues which we review next. The critical aspect is that operators have been covering large venues for almost as long as the mobile industry existed, but the trend is clearly aimed to deliver service to smaller venues. Today, dedicated in-venue service is available in many large venues such as stadiums, convention centers, airports, train and subway stations, and other large facilities that have high subscriber concentration. The challenge is to scale the service to cover smaller venues that include hotels, hospitals, and medium sized-industrial complexes. The proliferation of in-venue options is a response, or perhaps more accurately an anticipation, of the migration to provide service in smaller venues.

The Options

The options to provide wireless services to venues and buildings include:

Distributed Antenna Systems: Traditional DAS consists of passive RF devices and coaxial cables strung through a venue to distribute signals from a base station (Figure 3). Where losses are high, such as the case when the venue is large and the cables are long, or when the signals are split too often, bi-directional amplifiers are used to boost the signal strength. Passive DAS has a low cost-point but cannot scale effectively for large venues or multiple operators and frequencies. Active DAS solutions are best used to service such cases whereby the RF signals from the base station are converted to optical signals which are then transported over fiber a long distance to a remote radio where the reverse operation is done (Figure 4). Often, active DAS is combined with passive DAS for a hybrid deployment scenario. The extent of this practice depends on what the operator believes would work best in terms of project economics. When a hybrid deployment is considered, high RF-power remote radios are used to feed the passive network. Alternatively, the operator can consider a pure active deployment with low-power radios that are strategically located to meet the service level requirements for a venue. Active DAS systems are specifically targeted at large venues and can accommodate multiple technologies, frequency bands and operators with relative ease.

The active DAS worldwide market is valued at \$2 billion in 2013, up 2% over 2012, with a total of 1.2 million DAS nodes shipped. The global DAS market is forecast to grow at a 3% CAGR from 2013 to 2018, when it will top \$2.3 billion, and node unit shipments will pass the 2-million mark³. The overall DAS market including both equipment and services is estimated at \$4.4 billion in 2014 with forecasted growth to \$8 billion in 2019, of which a total of 60% will be on active DAS solutions⁴. North America remains the leading region for active DAS deployments followed by Asia and Europe.




Figure 3 Passive distributed antenna system.

³ Infonetics, "DAS market growth in N. America and Brazil offsets China slowdown," May 2014.

⁴ ABI, "In-Building Wireless Market Reaches \$8.5B in 2019," February 2014.




Figure 4 Active distributed antenna system.

Distributed Radio Systems (DRS): Are a relatively new breed of systems that extend the distributed base station architecture, a base station that features baseband processing module connected to a remote radio head through an optical interface (Figure 5). In the first type of DRS systems (Type 1), the baseband processing unit is connected through fiber to a low-power remote radio head (RRH) over an interface such as CPRI, which is most typical and exceeds OBSAI in adoption. An alternative system (Type 2) is the one recently introduced by Ericsson (DOT) and Huawei (LampSite) which uses an intermediary module to convert optical CPRI signals from the macro cell baseband modules into IF signals for distribution over CAT-type Ethernet cables to low RF-power remote radios. DRS provide the benefit of coordinating the operation among the low-power access nodes as well as between them with the overlay macro cell which can result in substantial gain in performance. DRS are also easier to plan, configure and manage compared to small cells, because a central baseband unit controls operations. However, DRS are limited in operating bandwidth to a few tens of MHz and in distance to a maximum of approximately 200 m, where CAT-type cables are used. The distance for fiber would reach up to a few kilometers. DRS are targeted at single operator deployments in medium-sized venues, especially ones where fiber is available.

The market for DRS is emergent at the time of writing this report with limited deployments as the solutions have recently been released on the market.




Figure 5 Distributed radio system of Type 1: low power remote radio, and Type 2: CPRI-IF conversion.

Small Cells: Small cells combine the baseband and radio frequency functions into one compact enclosure (Figure 6). They operate at different RF output power levels, ranging from a low of 0.2 W for indoor residential deployments to 5 W for outdoor carrier deployments.

Small cells can be deployed in two general network architectures. The first includes a gateway that performs certain management and security functions, which is typical for a residential and enterprise application. The second architecture is based on direct connectivity to the operator core network which is typical of carrier deployed small cells. Small cells are targeted at relatively small venues where DAS would be too expensive to deploy. Small cells are deployed typically by a single operator unlike DAS systems which are often shared by multiple operators.

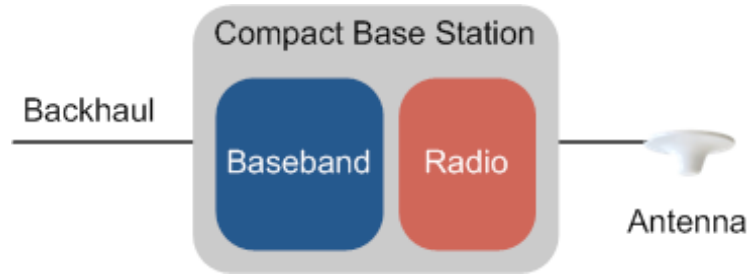


Figure 6 Small cell base station.

Wi-Fi: Wi-Fi is used extensively in the enterprise, SME and home as offload technology. Wi-Fi is also deployed in larger venues. Because Wi-Fi provides a low-cost point, it is believed that it will gain more popularity with operators to become an integral part of the radio access network. This objective is facilitated by recent technical developments such as the Hotspot 2.0 initiative which facilitates subscriber access to Wi-Fi based on the mobile SIM for authentication and security functions. However, Wi-Fi does not offer the same quality of service that LTE does, often because of poor planning or simply because of the high concentration of Wi-Fi access nodes. Hence, in deploying Wi-Fi, the network operator is faced with a classic trade-off between cost and quality. Nevertheless, Wi-Fi is a strong option for operators, and the technology has a rich roadmap that it is following, which will allow it not only provide better throughput performance, but more critically to better integrate with radio access networks.

Table 1 Comparative analysis of different in-venue wireless systems.

	Small Cells	DSR	DAS
Venue size	Small	Medium	Large
Management	Per module – controller/SON functions to reduce complexity	Per sector – through the macro base station. Follows general operator practices and systems	Per sector – through the macro base station. Follows general operator practices and systems
Potential for interference	High – requires coordination; SON functions can reduce complexity	Medium – requires planning. The distributed radios are coordinated among each other and with the overlay macro cell to reduce interference	Medium – requires planning. The remote DAS modules are extension of the base station sectors and coordination is possible to reduce interference

Distribution media	Fiber or copper	Mix of fiber and copper, or fiber only	Fiber
Potential for system sharing between MNOs	Low: depends on MNO attitude on sharing active infrastructure	Low: depends on MNO attitude on sharing active infrastructure	High: allows MNOs to install their own base stations which can be managed separately
Capacity capability	Supports single or dual frequencies with limits on number of users, typically up to ~60 for enterprise small cells, with new models reaching 200-400 users	Supports single or dual frequencies with higher limits on number of users than small cell. Limits per architecture and type of distribution network (e.g. copper)	Scalable with number of base station sectors installed. Cost as well as space requirements increase for large systems
MIMO	Inherent in the design and function of small cell	Inherent in the design of the RRH	Requires additional modules to support MIMO function that increases cost. New systems are addressing this with fully integrated modules, sometimes at the expense of lower RF power
CoMP capability	Low – the backhaul link would not provide sufficient capacity and jitter accuracy	High – intra site CoMP capability as distributed radios belong to a single sector of a multi-sectored base station	High – DAS extends base station sectors by operating at the antenna interface.

To round-up our review of in-venue communication options, it is important to mention Cloud RAN, an emerging technology that can play a significant and disruptive role once it matures over the next 3-5 years. Cloud RAN centralizes and virtualizes the baseband processing of the base station. This enables features such as coordinated multipoint (CoMP) where a mobile base station can communicate with multiple base stations simultaneously resulting in improved performance especially at the cell edge (Figure 7). From this perspective, Cloud RAN can be considered as an evolution of DRS to a higher level of integration and sophistication. In this report, we focus on delving deeper into the evolution of DAS to map its expected trends and development in the short and medium term (up to 5 years).




Figure 7 Coordinated multipoint.

Evolution of DAS Systems

A Perspective

It is perhaps useful to pause and review some of the history of DAS systems to frame the future evolution on what one can expect within the next few years. The roots of DAS are almost as old as the mobile industry. In the 1990's, operators started deploying what we now refer to as passive DAS as we introduced earlier. These systems consist of a network of coaxial feeder cable with taps to connect to antennas in different locations or alternatively a network of 'leaky feeder' cables which is a coaxial cable with gaps in its exterior conductor used to radiate energy (effectively slot antennas). Passive DAS performs relatively well for technologies such as GSM and for voice services running on 800 and 900 MHz where attenuation in coaxial cables is still relatively manageable, and the link budget would allow a distance up to a couple of hundred meters between the antenna and the base station. In cases where longer distances are required, bi-directional amplifiers (BDAs) are used to boost the signal strength in both the downlink and uplink paths (path imbalance is another major issue in passive DAS).

Passive DAS are susceptible to passive intermodulation (PIM) interference that result from mixing of different frequency bands, which increasingly became an issue the larger these systems got, with more frequency bands being added and more operators sharing a single system. As passive DAS systems struggled to meet the requirements in large multi-operator venues, active DAS systems emerged as a solution. Passive DAS does not support fault management capability (alarms) nor does it allow power management and control capability at the antenna level. Yet, passive DAS systems remain a low-cost option that is used whenever the size of the venue supports such deployment. Passive DAS as the name implies does not include any active modules, with the exception of BDAs which are simple, low-cost devices. Once installed, passive DAS can generally operate for many years into the future, especially inside buildings where the environment is controlled.

The Present

Active DAS systems evolved to solve many of the limitations of passive DAS. Active DAS provides long reach and better protection against PIM by converting RF through an intermediate frequency (IF) down-conversion stage to optical signals in a master or DAS head which are then transported over fiber optical cable to a remote location where the reverse is accomplished. A remote unit converts the optical into RF signals that are amplified and transmitted. While the concept is relatively straight forward, the implementation and design of active DAS systems is a basis of differentiation between vendors. The DAS systems on the market today were primarily designed to cater to the established technologies and frequency bands used by operators: GSM, CDMA/EV-DO, and 3G/HSPA running in 800/900, 1800/1900, and 2100 MHz bands. In fact, some systems are limited to a certain technology and band which is becoming a challenge for the current operating environment as operators today have increased their spectrum holdings and operate multiple technologies.

Active DAS provides the network operator with management and control capabilities including fault management. Active DAS systems connect to the base station through a Point of Interface (PoI) which consist of RF signal shaping modules (splitters, duplexers and multiplexers, couplers, attenuators, matched load, etc.) to condition the output signal from the base station which is generally at high RF power (order of Watts), to a level that is suitable for input into the DAS (order of milli-Watts). The PoI is one of the cost drivers for DAS especially for large systems that include many operators, frequency bands and carriers. Moreover, the bulk of base station RF output power is dissipated in a matched load which is inefficient use of energy. PoI modules also consume space which can be limited in many venues.

Considering there are three main elements to active DAS systems (PoI, master head and remote unit), active DAS differentiate by how these elements are designed and how they work together to form a complete system. The characteristics and the way these building blocks are assembled and interconnected to deliver on the coverage and capacity objectives for a certain venue (small or large) result in different cost structure which favors one vendor solution for a certain deployment over another. In other words, one aspect to DAS systems is that there is no single universal solution that is superior for all use cases – there is ample opportunity to differentiate and to focus on specific target markets and applications. This is evident by the path that vendors have taken in designing their systems. Here, we focus on two aspects: the optical distribution system and the remote radio.

Optical Distribution: There are fundamentally two modes for optical signal transmission over fiber cable: analog and digital. The majority of DAS solutions on the market today are based on analog modulation of optical signals by RF signals. Typically two fiber optical cables are required to connect the master unit with the remote unit with one for each direction, the downlink and the uplink. The second mode is digital modulation of optical signals. In this case, some systems use two different optical wavelengths and combine the downlink and uplink signals on one fiber optical strand. Digital systems, whose presence on the market is increasing, have the capability to deliver longer range than analog systems because of better optical power budget. This has the potential to allow new business models centered on base station hosting. Digital DAS also allow the operator to switch signals from one remote radio to another which allows them to serve different areas using the same baseband resources, thus reducing cost of deployment (Figure 8).

They also provide high flexibility in providing different deployment topologies that optimize the design of the distribution network for lower cost for example where systems can be deployed in a star, chain, loop, or hybrid topology. Analog technology, on the other hand, is widely available at relatively low cost-points and can be used effectively in scaled down DAS systems into smaller venues where sensitivity to cost increases. Moreover, some analog systems can support very wide bandwidth which allows supporting greater number of RF carriers in the optical distribution system. Specific examples of digital DAS include that of TE Connectivity and Dali Wireless. Axell Wireless and Commscope also announced digital products recently in a shift from their traditional analog systems. In all, we see a trend to deploy digital systems in larger venues and in campuses where range, capacity switching and other features combine for an effective business case. On the other hand, analog systems can scale faster in cost to serve smaller venues.

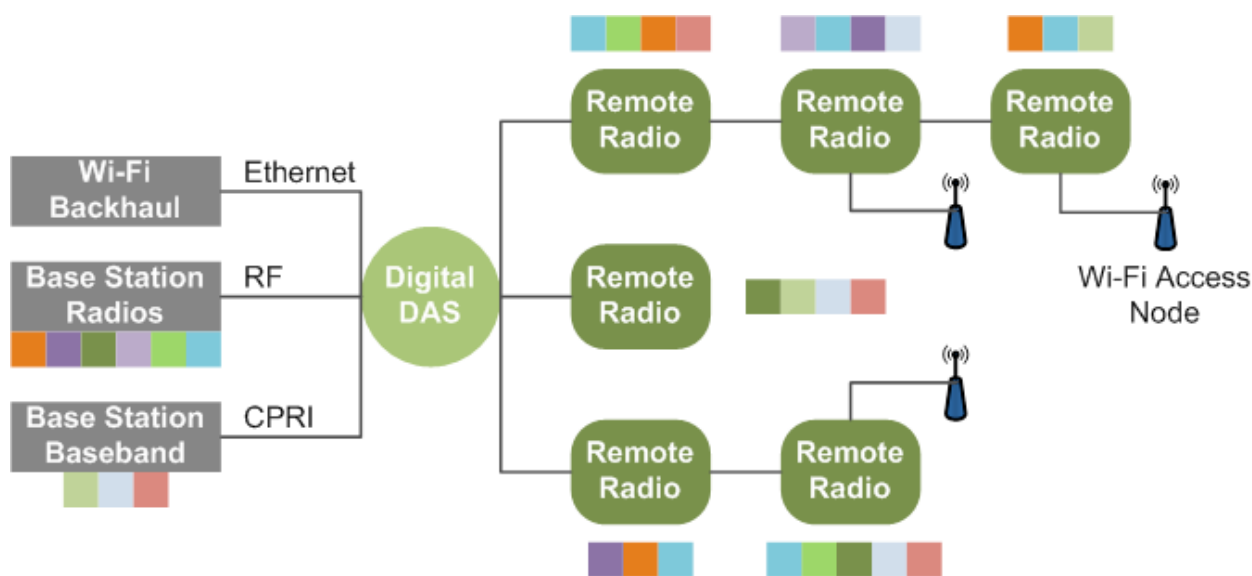


Figure 8 Digital DAS enables switching RF signals from the base stations to any remote radio and enables greater integration with the base station and Wi-Fi access nodes.

Remote Radio: Two of the main characteristics of remote radio are the bandwidth and output power. Here again, vendors have differentiated their solutions. Most remote radios on the market accommodate multiple frequency bands in different types of enclosures. Remote radios come in different RF power outputs: sub 1 W (low), between 1-4 W (medium) and 4-20 W (high). High power radios can be shared by greater number of operators because the power is divided among the different users of the system. They can also be used to feed passive DAS networks. On the other hand, low power radios have a relatively small size and can be easier to deploy and used in greater quantity to provide uniform coverage and performance. Recent DAS have implemented digital pre-distortion and crest factor reduction techniques to improve the performance and reduce power consumption, a trend that will continue to spread. Aside from output power, the bandwidth capability is another critical factor. Wide bandwidth allows greater flexibility in spares, inventory management, and flexibility of future growth. However, higher bandwidth typically comes at a price or lower RF power output. Here, we point to a specific example of Zinwave's unique wideband radios that support all wireless frequencies between 700 MHz – 2700 MHz in a single low power module.

The above exposition of active DAS systems demonstrates multiple approaches taken by vendors. Each solution provides distinct advantages, which makes it imperative to consider different options for a specific venue that accommodates service design objectives.

Evolving Trends

The evolution of DAS has to factor the evolution of market requirements such as scalability of DAS to smaller venues. This requires a reduction in cost, the simplification of installation, availability for deployment and management by third parties, as well as improvements to size, form factor and aesthetics. The challenge lies in that these requirements are accompanied with the need to support greater numbers of frequency bands and different technologies (e.g. HSPA, FD-LTE, and TD-LTE). The downward evolution towards smaller venues does not exclude continued evolution to provide higher cost efficiency for large venues. In fact, this is where digital-based DAS systems provide much value. Therefore, the evolutionary trends are two pronged with the first focused on the downward trend into smaller venues, particularly in developed economies and markets, and would have wide market implications mainly because many of the venues are green-fields with no current service. This is an area that will place DAS in competition with DRS and small cells. The second is focused on achieving greater cost efficiency for large venues, which opens up new markets in emerging markets as well as new applications in the developed economies (e.g. base station hosting service) (Figure 9). In this sense, we expect the DAS market to branch further as more use case scenarios become possible.

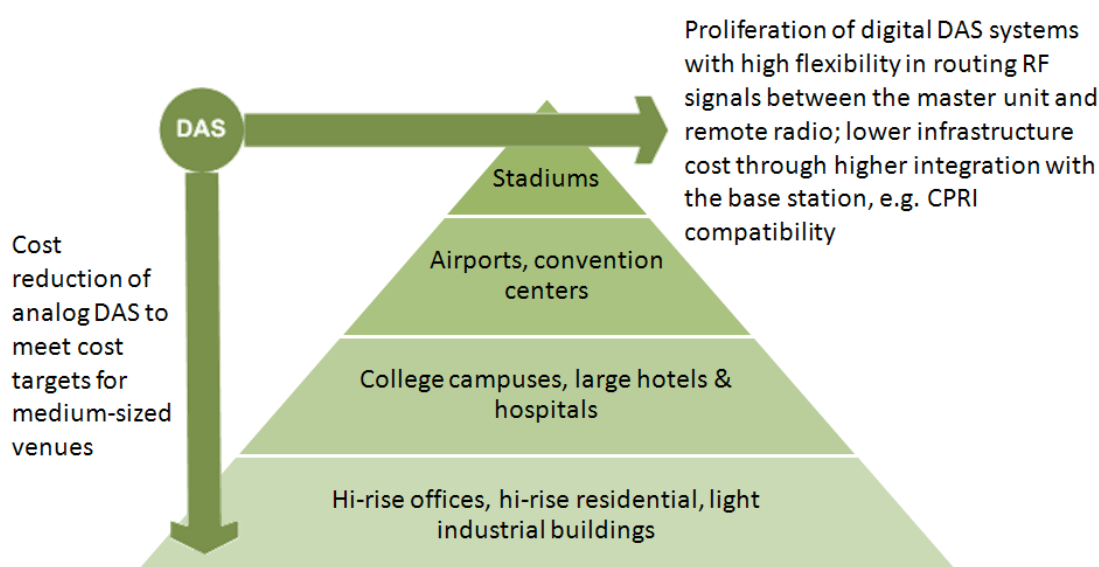


Figure 9 Evolution of DAS to reduce cost for large venues and to scale service into smaller venues.

CPRI Integration: A notable emerging trend is to extend DAS to support CPRI standard directly from the baseband unit. This has an advantage in reducing the cost of deployment as it eliminates the need for a radio head at the base station which saves significant capital expense. The BTS radio, which typically accounts for as much as 50% of the base station hardware cost, and the PoI associated with the DAS, and operational expense related to energy consumption are all saved. There are, however, important consequences to implementing this approach. The first is that base station vendors control the management and control layer of the CPRI interface. The DAS vendor is required to collaborate with the base station vendor to realize full and seamless integration.

The second consequence is that CPRI consumes very wide bandwidth. For example, a single 20 MHz LTE channel with 2x2 MIMO support requires 2.5 Gbps line rate. This would quickly use up the capacity available on a fiber cable and can necessitate the use of WDM to combine multiple CPRI signal streams on the fiber cable, or alternatively use more cables. Nevertheless, CPRI integration is a significant feature that positions DAS close to DRS and creates direct competition between the two approaches. Hence, we see Alcatel Lucent opting to integrate with TE Connectivity to provide a function similar in concept to Ericsson DOT system. Another company stating CPRI compatibility is Dali Wireless. Both of these companies have developed digital DAS which is amenable to carry CPRI IQ data.

Wi-Fi/Ethernet Integration: Wi-Fi is widely available indoors in the enterprise as well as for public access. Integration between DAS and Wi-Fi is therefore a logical step, whereby the fiber infrastructure used by the DAS system can be leveraged to carry Wi-Fi backhaul traffic to a central location in the venue. Digital DAS are increasingly equipped with one or more Ethernet ports at the master and remote radio unit to multiplex Wi-Fi Ethernet backhaul signals. Some systems support only 100 Mbps, which is relatively small, but newer systems support 1 Gbps interface which provides greater capability to support Wi-Fi access node. The result is lower cost of providing wireless coverage and Wi-Fi data service inside a venue.

MIMO Support: MIMO presents a critical implementation challenge to DAS and has exposed a current weakness, where implementation requires doubling the entire distribution system which essentially doubles the cost of the deployment. MIMO requires a completely separate RF-optical conversion module, a remote radio and the fiber connecting them. This is because MIMO spatial multiplexing consists of different information bit streams transmitted at the same frequency. In DAS, the two streams are required to be separated for processing and to eliminate interference between the two streams, hence, the effective doubling of DAS hardware requirements. MIMO is a feature of LTE so while LTE networks are still lightly loaded today the pressure to deploy MIMO is not urgent, giving vendors some time to develop cost effective solutions. Nevertheless, MNOs have favored MIMO deployments in venues and consequently provisioned for MIMO in DAS deployments. With future wireless networks relying on a greater order of MIMO to achieve capacity, the challenge to support MIMO in DAS will increase proportionally. This will be a key area where DAS, DRS and small cells will compete and differentiate.

Multimode Support: Today, we can find several wireless technologies operating in the market: GSM, 3G/HSPA and LTE (FDD and TDD). Additionally, the evolution of LTE comprises different operating modes such as carrier aggregation, which incorporates an additional carrier as a supplementary channel to augment the downlink path. Today, most DAS solutions on the market are limited in their ability to support the TDD mode. This presents a challenge to sharing the DAS with TDD operators. Multimode FDD/TDD support is another cause for DAS evolution.

Alongside the four developments above, there are four evolutionary trends for DAS to follow:

Frequency Axis: DAS will have to evolve to support wider channel bandwidth and a mix of different frequency bands in conjunction to increases spectral holding of MNOs. This will allow multiple operators to share a system, resulting in greater cost efficiency. The capabilities of DAS to cost effectively support a varied mix of frequencies is a key differentiation point especially valued by the DAS operators.

Power Axis: DAS will evolve to support different variations of radios with multiple output powers. Specifically, medium power modules in the range of 1 W would cater well to the smaller buildings while at the same time allowing multiple operators to share the system as the case requires. This power category of radios would see high growth.

Integration axis: To reduce the cost of DAS deployments in medium-sized venues, it is possible to use small cells as feeders into the DAS. Integration of small cells and DAS into a single operating system provides both coverage and capacity at reduced price, by eliminating the macro cell and reducing PoI requirements because small cells operate at lower power. For this to succeed, the combined solution would use high-capacity small cells (e.g. 200-400 active subscribers) that have recently become available on the market. Another aspect of integration pertains to the capability on the optical distribution network where greater use of WDM solutions is anticipated to increase the utility of DAS.

Deployment and operation axis: DAS projects are typically large and require coordination between different entities to bring about a fully deployed and operationally effective system. In scaling to medium sized venues, ease of deployment will take on added importance, as often it will be third parties who would install DAS. Means to ease deployment can take different directions, such as reduction in space required for the DAS, auto-calibration for near plug-and-play installation, capability to use different media for transport of signals between the master and remote units. In addition to this there is fiber and other features that help make the deployment process simpler and more cost effective, such as the use of single fiber for both downlink and uplink paths. Furthermore, the systems need to be managed in a straightforward manner at an independent operator level. Greater functionality in software will be a key differentiator in DAS that will gain prominence.

The above trends would combine to extend the capabilities of DAS to render them simpler to deploy and easier to maintain. The main appeal for DAS has been the ability to provide a single point of inter-connect to the base station which, at least in the United States, has provided a demarcation point between the mobile network operator and a third party who designs, deploys and maintains the system. This model will slowly make its way to other regions in the world and prove to be a catalyst for continued growth in DAS.

Enabling New Business Models & Applications

The evolution of the wireless base station architecture to include small cells and Cloud RAN, in addition DRSs, increases the competitive pressures on the DAS vendors as more options are available to the MNOs for in-venue service than ever before. Yet, there are distinct features and advantages to DAS that would keep it as a viable option for many venues and certain types of applications. One of the benefits is that DAS can be easily shared by multiple operators which reduces the capital and operational costs. In contrast, small cells, Cloud RAN or DRS require sharing of active infrastructure. MNOs in many markets, especially those where ARPU is relatively high, refrain from adopting this as operators continue with the strategy of differentiation based on network performance. Another relevant feature of DAS is that they can be deployed by a third party who provides MNOs with a single connection point to the base station. MNOs favor demarcation points where responsibility for service and support can be clearly identified.

With this in perspective, the advantage of DAS lies in the business model and applications that it

enables. These are largely driven by an operator attitude towards system sharing and third-party engagements. The business model factors heavily into the resulting cost of deployment, which has played favorably for DAS in large venues. As such DAS will evolve to accommodate greater integration with the base station as will be required in the future, especially for large venues. At the same time, the evolution of alternatives will continue to create more tension among all these technologies as each technology progresses along its development path. The success of one over another would largely depend not only on which is better able to accommodate the preferred business model, but also on what a technology provides in new business models and applications, which are bound to vary among different regions and markets.

As an example, the new generation of digital DAS enables new business models centered on base station hosting. The high optical power budget of digital DAS allows aggregation of base station baseband in a central fiber office removed by tens of kilometers from the remote radios. This application has been used to some extent in outdoor DAS deployments, but the new systems would provide greater benefits. Such as where base station hosting can be coupled with capacity switching to serve moving traffic hotspots. This conserves base station baseband resources as these resources would be pooled and assigned dynamically to hotspots as required. In effect, the benefits are similar to what Cloud RAN provides, as it is no longer required to provision capacity for the peak value required for every location. As a practical example, DAS can be used to provide service over long stretches of railway tracks with minimal baseband resources that are switched from one remote transceiver to another as a train passes through its coverage area. This application provides a railway company or a subway operator an opportunity for additional revenues should it decide to deploy such a service. In a correlated model, a fixed access operator with fiber assets can provide base station hosting service in its fiber centers and use its access to commercial buildings to enable the MNOs to serve these buildings using its already deployed fiber. Note that such a case can also be implemented with small cells, provided operators are more amenable to sharing infrastructure.

Conclusions

There is heightened attention on in-venue communication systems as a means to improve wireless services that are taken for granted by subscribers expecting service anywhere, anytime. This attention is augmented by the need of MNOs to offload congested macro cells by eliminating traffic hotspots through the lowest cost alternative, leading to a convergence of objectives that has combined to stimulate growth of DAS solutions. While other solutions that include DRS and small cells are alternatives for in-venue solutions, DAS was and continues to be the workhorse, mainly because the business model it provides has been amenable to operators. Starting with deployments in the largest of venues, the evolution of DAS is expected to continue along two paths, one leading to lower-cost deployments and the other realizing DAS economics for smaller venues. The market for DAS is expected to continue to grow in the absence of a consensus by operators on infrastructure sharing. The evolution of DAS would center on enhancements of digital distribution technology that allows higher cost efficiency and integration with wireless base stations to reduce total cost of ownership, in addition to the introduction of low-cost analog based solutions with greater flexibility to meet the requirements for relatively small venues. The viability of DAS in the future would hinge on the new applications and business models it can enable. In this paper, we provide examples of applications that new generation of DAS enable by leveraging the flexibility of the digital architecture for cost effective new deployment scenarios.

Acronyms

3G	Third generation
ARPU	Average revenue per user
ARQ	Adaptive repeat request
BDA	Bi-directional amplifier
BTS	Base transceiver station
CDMA	Code division multiple access
CoMP	Coordinated multipoint
CPRI	Common public radio interface
DAS	Distributed antenna system
DRS	Distributed radio system
EV-DO	Evolution - data optimized
FD	Frequency duplex
FDD	Frequency-division duplex
GSM	Global System for Mobile Communications
HSPA	High speed packet access
IF	Intermediate frequency
IQ	In-phase and quadrature
LTE	Long Term Evolution
MIMO	Multiple input multiple output
MNO	Mobile network operator
PIM	Passive intermodulation
PoI	Point of interface
QAM	Quadrature amplitude modulation
RAN	Radio access network
RF	Radio frequency
SIM	Subscriber identity module
TD	Time duplex
TDD	Time-division duplex
TE	Tyco Electronics
WDM	Wave division multiplex

XOONA PARTNERS

Home based Healthcare & Opportunities for Mobile Operators

Ananda Sen Gupta,
Dr. Riad Hartani, Richard Jeffares

February 2014

Internet Technology in the Health Care Eco-system: Rationale

The World Health Organizations global status report on non-communicable diseases 2010 outlined non-communicable diseases (NCDs) as the leading global cause of death, responsible for more deaths than all other root causes combined, and NCD's strike hardest at the world's low and middle-income populations. These diseases have reached epidemic proportions, yet they could be significantly reduced, with millions of lives saved and untold suffering avoided, through early detection, timely treatments and reduction of their risk factors.

Of the 57 million deaths that occurred globally in 2008, almost two thirds were due to NCDs, comprising mainly cardiovascular diseases, cancers, diabetes and chronic lung disease. The combined burden of these diseases is rising fastest among lower-income countries, populations and communities, where they impose large, avoidable costs in human, social and economic terms. About one fourth of global NCD-related deaths take place before the age of 60.

Given the fact that the numbers of Doctors being educated and entering the workforce is never going to catch up to the speed at which the disease burden is increasing globally, there is a need to make fundamental changes to the process of offering healthcare and these changes need to make it more agile, efficient and robust.

Whilst an opportunity exists for the introduction of wireless and portable technology solutions, it is important to acknowledge that healthcare has traditionally been a closely bound and guarded segment which was largely limited to human analysis and intervention, with technology applications primarily used in specialized diagnoses. The problem is aggravated by the fact that though scientific knowledge is present in the area of human anatomy, data is typically not available in real-time regarding the effects of continuously changing environmental factors on health conditions. This has made it quite difficult to predict illnesses for an individual. On the other hand, for chronic illnesses, there is now enough evidence on why and how health conditions deteriorate due to poor choice of lifestyle.

This paper addresses the opportunities emerging for mobile network operators and cloud solution providers to take advantage of the next evolution in health care, and describes a home based health care application use-case designed and deployed to illustrate such opportunities. Specifically, it addresses the potential synergetic models developed between health care providers and mobile operators.

Mobile Health as an opportunity to Mobile Operators

Mobile operators globally are in a phase where two options are put in front of them: either to optimize their networks to becoming a mobile broadband path, with no or little plans to share a piece of the revenues derived by the Over the Top (OTT) players, or to position their network, selectively, within the overall OTT value chain, to share a piece of the revenue streams. This is also the case in the context of mobile health, where some operators, have been, and are still, working on defining their own approach to this market, now that mobile devices penetration is high, smartphones/tablets offer screens large enough for advertising and revenue streams off mobile health care are seen as a good alternative to declining revenues in traditional voice services.

In the mobile operators' favor is the existing subscriber relationship, where location aware applications can couple with health-monitoring data that is streamed real time to the mobile

health professional or big data repository. A number of challenges still remain, as mobile health has never been in an Operators DNA historically and significant transformation is required to support small payload M2M styled traffic models internally. Another fact is that many OTT's have already entered this market aggressively with smart phone and tablet applications, making it difficult for newer entrants to clearly differentiate their value upon market entry.

Mobile in Home Enabled Health: Foreseen Evolution

There is general acceptance now that preventing or delaying the shift of patients to acute- or long-term-care settings, is of enormous value that can be seamlessly provided by technology enabled home care. This is because the direct costs associated with any other care facility outside of the patient's home are significantly higher.

While, any technology used in home care cannot address all the potential factors underlying such shifts—for example, an accident. Health professionals agree that the medical conditions that can be addressed successfully by technology-enabled home care are as follows:

- Chronic conditions – conditions that persist for years rather than for a short while.
- Conditions that can be prevented or addressed by protocols, i.e. repeatable and standardized set of instructions that can be executed by non-physicians as well.
- Conditions which do not require round-the-clock attention or intense human monitoring.

Key Success Factors

The key factors of success of this model are as follows:

- 1. Clear and significant impact:** A home health care model and technologies must provide information that can be effectively used to affect the patient's overall course of disease progression and plan appropriate interventions. For example, monitoring the weight of a patient with congestive heart failure can provide early warning to the clinician to imminent worsening of patient's condition. Again, by analyzing various vitas data on a regular basis can provide a good understanding of a hypertensive person's health while they are on medication.
- 2. Timely and Actionable information:** Simply observing parameters or creating a health alert based on the data collected using the home care technology is not meaningful enough. There has to be a way to take appropriate action, be it through a caregiver, nurse or emergency support service, when such an intervention is deemed necessary. For example, an emergency intervention may be required in case of a sudden weight gain in a congestive heart failure patient, instead of simply providing a weight gain chart on the screen.
- 3. Closed loop approach:** A home based health care solution (which will be a combination of team, products and processes) must have a closed feedback loop so as to measure progress against the goals that have been set, and understand if actions and treatments have been effective or not. Processes and data collection process has to be seamless, so that the feedback does not get overlooked in any way. To complete a closed loop, the processes and health support team have to be fully involved to take timely action based on the measurements.

4. Easy to use and automated as far as possible: The home health care technology must be simple to use and appreciate by the users. The automated wireless blood pressure measurement device used at home without major technical understanding is way more easily usable than a fixed blood pressure kiosk at a pharmacy. Also, any technology has to be designed for a large population, and not for controlled trial population scenarios.
5. Recurring readings: The technology must be used to take regular and frequent readings. The daily measurement of body weight on an electronic scale by congestive-heart-failure patients is repeatable. Any product that is only required to be used intermittently is not valuable for home use.
6. Clear financial benefits: The return on investment (RoI) for the implementation of home care technology must be clear to patients. Typical Personal health record software for patients, for example, could never become popular because users need to enter a great deal of information manually in return for ambiguous benefits. However, if there is an organization which helps makes sense of the collected data and then provides distilled information to the Doctors and the patients as well as their caregivers, the value will be clear and direct.
7. A clear connect between payers and providers: Health care service providers such as private hospitals can feel left out in the process of home based healthcare, as they may consider this a loss of revenue. Smaller companies may start playing this role and fill this gap. On the other hand, at a Governmental level where the Government is both the payer for as well as the provider of the services, the overall cost burden on the health program will go down steadily with effective implementation of home based health care.

Mobile Operators: How to approach the mobile health opportunity

Multiple options are being considered in terms of how to approach the mobile health market. They are described below.

Option 1: Mobile networks directly acquiring mobile health players to build a direct presence in this space.

This is the case of the largest mobile networks, with an aggressive push towards mobile health where a dedicated and scalable ecosystem needs to be created. One approach is to do this through the acquisition of relevant players or acquiring a significant commercial position, which in turn provides the growth option of building a business upon these new technologies.

Option 2: Mobile networks partnering with mobile health players to build a direct presence in this space.

This is the alternative approach that some mobile operators have considered, as a strategy to approach the mobile health market. In these cases, the platforms are owned distributed and managed by the partners, but through a well-defined partnership model with the mobile operator.

Option 3: Mobile operators build their own mobile health platforms to compete directly with mobile health centric players.

This is the case where operators have gone into designing and implementing their own mobile health solutions and underlying ecosystem. This is still in early stages of development, but in some cases, operators have been working on sharing common co-developed M2M platforms to address the fragmentation problem and increasing the size of the customer base and having it approach the addressable size, as seen per an OTT. This is specifically the case of small mobile operators who would need to join efforts to get to a sizable customer base.

As a complement to such models, some operators are looking at having their own mobile health integration within their branded App Stores, as a way to counter initiatives of larger OTT players.

Option 4: Mobile operators focused on defining new business models leveraging mobile health without directly managing the mobile health eco-system.

In this case a number of mobile operators have done so in conjunction with one of the 3 options above. In most scenarios, this is built upon the existing operations process of mobile operators, such as performing content re-formatting based on screen size and/or formatting, augmenting billing models to accommodate mobile health information insertion models, augmenting their marketing campaigns with mobile health related information at retail point of sale, leveraging data warehouse information to be exposed to the mobile health eco-system running on top of the network, and to lastly integrate mobile health with content distribution networks in-house.

It is worth noting that within each of these various models, mobile operators aim at inserting themselves into the mobile health value chain from different angles, based on a strategy that is optimal to them. One should note that although various models are being considered, various challenges still exist for mobile operators at a regulatory level.


Mobile health is as yet not fully defined by the various governments in some parts of the world, where it will simply remain an enabler, whilst in more evolved markets, it will become a key mode of vital signs monitoring and underlying healthcare delivery .

Mobile operator management teams have little experience dealing with the various actors of the mobile health eco-system, and must address various privacy considerations in their locale, as well as the customer expectation management challenges that mobile health potentially introduce.

Case Study: Trackmybeat Healthcare

Trackmybeat Healthcare has created a solution that allows simple, easy to use and familiar medical diagnostic devices to collect key medical parameters from the patient-home and send the data real-time to a central data store through a mobile App (application running on a mobile phone) automatically.

Detailed analysis is then conducted on the collected data continuously and, based upon the resulting information a Doctor can plan a suitable treatment revision or direct clinical intervention.



Here are two potential models of how a Mobile Operator is planning to integrate this into their service portfolio.

Model 1 (following Option 3 as above):

A Mobile Operator is considering adopting the Trackmybeat solution as part of their expanding OTT App Store for Health and Wellness, and offer the remote health data collection service and analytics results to both individuals, as well as Health Service Providers such as hospitals, and state government health departments.

Model 2 (following Option 4 as above):

A Mobile Operator plans to offer Data Centre services to Health Service Providers, where they host the Hospital Information Services (HIS) solutions. With that, they are planning to partner with Trackmybeat, so that they can offer remote data collection services to the Health Service Providers, as well as analytics of the data to feed to the HIS database of the Health Service Providers, for better informed clinical decisions.

Conclusion

It is imperative for Mobile Operators to understand the changing healthcare service landscape and adopt suitable business models around it. As usual, one size will not fit all. This is due to the fact that in different markets, there may be different payers for healthcare services, as well as different regulatory environments. This will affect the ability of the mobile operator to offer specific services within the space.

Hence the mobile operators will choose from a variety of strategic options, which range from acquiring mobile health solution providers to offering infrastructure support that is tailored to mobile health service providers.

A real world case study has now been developed and commercially deployed by the Xona team, with high-level lessons learned outlined within this paper.

XONNA PARTNERS

Spectrum and Network Sharing Models Trends & Business Impacts

Dr. Riad Hartani, Frank Royal
Andrew Murton, Richard Jeffares

January 2014

Synopsys

Forward-looking access service providers have amassed considerable fiber optical assets complemented by Wi-Fi services. These operators are now considering the next stage of revenue generation and growth. This paper discussed the synergy with Mobile Network Operators (MNOs) in light of ongoing business model and technology developments, with a focus on evolving network and spectrum sharing trends. It also addresses the direct impact on service providers' business models, and concludes on the likely emergence of large-scale Internet and cloud-centric virtual operators.

Network Sharing Models – Situational Overview

There are two angles to the resource-sharing models: the first angle relates to passive infrastructure sharing which is being pursued throughout the world in various forms (e.g. tower sharing) and active infrastructure sharing, and through an extension of it, of spectrum resources. It's the latter angle that is now gaining direct attention where various models under experimentation.

Over the last few years, various forward-looking operators, and specifically fixed-line operators, have taken the lead in building high-speed fiber access networks (FTTx: Fiber to the Premise/Home/Curve). With such achievements, they have, without explicit planning, put together the initial building blocks for global leadership in optimized neutral host and infrastructure sharing service. In fact, as 3G and 4G networks got deployed, requirements for high-speed backhaul grew, which provided some of these operators with a unique opportunity to leverage their fiber infrastructure for this purpose, mostly as wholesale backhaul capacity to mobile network operators. The rapid increase in 3G and 4G capacity requirements driven by the bandwidth requirements of over-the-top applications led to a fast-growing need of complementary technologies to accommodate the growth in demand for capacity. This in turn provided these operators with the opportunity to augment their fiber networks with Wi-Fi rollouts, and leverage Wi-Fi assets as complementary building blocks for their neutral host infrastructure sharing plans through Wi-Fi wholesale and offload offerings. With this, both the fiber and Wi-Fi infrastructures would form the backbone of these operators wholesale and infrastructure sharing strategy.

Active Infrastructure Sharing – Market and Technology Trends

Given this development, the question converges on what additional technology deployment strategies would be required to re-enforce and augment the infrastructure sharing model? Few propositions could be positioned, but the most immediate and relevant would be a direct complement to the backhaul and Wi-Fi plays that would simultaneously address the common customer base of both Wi-Fi and backhaul services (i.e. MNOs and enterprise / business venues), provide a direct competitive advantage against potential competitors, and solves some immediate problems faced by this specific customer base.

In analyzing the various arguments, the following is emphasized:

(a) The most urgent concern of MNOs is to optimize capex and opex while they augment their coverage and capacity requirements. DAS and small cell buildouts are specific areas where this concern is acute, and hence, MNOs are receptive to business models that would allow them to build such complementary networks while keeping their costs in check.

(b) The competitive MNO environment in different telecom markets and the stringent requirements of end users, be they business venues or the customers of such business venues, is forcing MNOs to act promptly on their network coverage and capacity upgrades that adds a time constraint dimension to the capex/opex considerations.

(c) The trend of mobile operators in some markets having to increasingly compete on services rather than coverage, is forcing them to put their energy into the services layer, which provides them an incentive to share more network resources to meet coverage objectives.

(d) In some select telecom markets (example Southeast Asia, Africa, Middle East), some lead operators are in a unique position where, as a non-mobile operators (so far), they are perceived to be neutral and not a threat by the MNOs which is conducive for strategic partnerships.

(e) The architecture of the wireless base stations has evolved to a split architecture that separates the baseband processing from the radio module. Many operators have already or are in the process of migrating to this new base station architecture, which requires fiber connectivity between the baseband module and the remote radio head. The fiber connectivity is referred to as ‘fronthaul’ and is seen as complementary in function to backhaul that connects the base station baseband module to the core network. This provides a unique opportunity for select network access operators with substantial fiber deployments to provide fronthaul as a service expanding on existing backhaul business with the MNOs.

Pushing Ahead with DAS and Small Cells

Focusing on DAS and small cells technologies with the above in mind, two fundamental questions need to be considered:

(1) What strategy to consider in successfully implementing a DAS and small cell infrastructure-sharing business model?

Our detailed analysis of the vendors and their offering in this space, technology readiness, MNO readiness, acceptance and leverage in select markets (Southeast Asia and Middle East) concludes that a shared active DAS deployment model would be the first step to consider mostly because sharing (specifically for passive DAS, and to a large extent for active DAS) is already a common practice between MNOs. Upgrades from passive to active DAS systems are becoming required with the roll out of LTE, particularly as LTE offers high data rates at modulation levels that require good signal quality which passive systems will be challenged to provide not to mention the opening of new frequency bands in 2300 and 2600 MHz that stresses the capability and performance of passive DAS systems. Such developments require fiber connectivity and ultimately provide network access operators with leverage in commercial venues and business relationships.

In parallel, a small cell sharing strategy (including the small cell / Wi-Fi combo solutions) would be built initially on the basis of optimized shared backhaul to small cell sites, and over time evolve to shared small cells when the technology is ready (multi-frequency/channels, virtualized management, etc.) and sufficiently mature to be deployed in a multi macro/small-cells vendor environment where MNOs allow third-party management of the small cells network. As such, priority is currently on active DAS shared deployment first with the building blocks of a small cells sharing model to be put in place over time (backhaul/fronthaul then small cells).

This strategy is enforced by difference in applications between DAS and small cells where the business case for DAS is more efficient than small cells in large venues while small cells are more efficient for small venues.

(2) Given the strategic investment by some of the operators we have analyzed in Wi-Fi, how would such shared active DAS deployment complement the overall plan, and what else could be done to re-enforce it?

Today's Wi-Fi and DAS/small cell networks are distinct, and could play complementary or competitive roles based on how they are positioned. Most MNOs see Wi-Fi and DAS/small cells as complementary, addressing different traffic profiles, usage behaviors, geographical fit, etc. In other words, they are likely to co-exist for the foreseeable future to address complementary needs. Synergies do however exist between these technologies where the possibilities include: leverage of a common management/authentication backend in MNOs' networks, leverage of common user billing platforms, and leverage of similar VAS (specifically if Wi-Fi traffic is backhauled to the network core). At the same time, these technologies re-enforce each other when it comes to new customer acquisition and/or customer retention. With such MNOs having lead on the Wi-Fi angle, a lot of what is already done with Wi-Fi can be leveraged as per the above, from common fiber and backhaul infrastructure already built by these forward-looking operators for their Wi-Fi, to common backend/billing/management, to interaction with common customers/venues that would benefit from the complementarities of Wi-Fi and the DAS/small cell infrastructure. As such, having Wi-Fi and the underlying infrastructure in place highly increases the value proposition of these operators in positioning a sharing model with Wi-Fi and DAS continuously re-enforcing each others in terms of value to the MNOs as the shared infrastructure is built.

New Opportunities Beckoning – Towards Cloud RAN

We have already mentioned that the evolution of the base station to a split architecture introduced the concept of fronthaul, which is the connectivity between the baseband and radio modules. While this can be considered a complementary concept to backhaul, significant differences exist which are driven by the technical requirements. Fronthaul requires an order of magnitude greater capacity than backhaul and is subject to stringent requirements for other technical parameters like delay and jitter. MNOs looking to maximize performance have an option to deploy Cloud RAN architecture in the future where centralized baseband processing drives a number of remote radio heads. The remote radio heads can be deployed in macro cell configuration or in small cell configuration. In both cases superior performance can be achieved over traditional distributed architecture (average 20% on uplink and 5-15% on downlink). To realize these gains, the business case for dark fiber for fronthaul needs to be sufficiently attractive. This is another area where forward-looking access operators can aim at. In our studies of the market, we developed regional business cases that flush out the important parameters for the success of this idea.

Cloud RAN architecture aims to decouple the base station software from the hardware platform which is reduced to COTS servers augmented by processing engines for computationally intensive physical layer operations. In this, Cloud RAN may well open new schemes of infrastructure sharing and/or neutral hosting models especially in markets where the fiber operator is neutral or is a

MNO that does not consider competing on network quality and performance more advantageous than competing on price or service. This case leads us to believe that even more creative business models may develop to operate the wireless network such as the cloud-centric virtual operator discussed below.

In the meantime, access service providers do not have to wait for the full evolution of Cloud RAN. Digital active DAS allows MNOs a similar deployment model and capability to extend coverage into hard to reach areas for base station deployments. In this scenario, the base stations are located at the fiber central office with long runs of fiber (typically < 10 km) to remote radios.

Spectrum Sharing and Shared Spectrum

Sharing spectrum assets between operators has proven to be contentious in many markets in part due to operators' own competitive behavior, and in other part due to regulatory rules. Yet, we do see many examples where MNOs came to the understanding that service and revenue generation trumps capital and operational expenditures necessary to maintain competitive edge in performance and quality which is not sustainable in the long run due to the nature of wireless signal propagation, coverage performance, and interference management that are critical for capacity. In other words, there is a plateau in service quality and diminishing returns to expenditure on network quality. Additionally, MNOs in markets where ARPUs cannot sustain continuous development for high level of service performance have taken the pragmatic lead to share spectrum and the radio access network to provide a better service than otherwise would be possible.

Today, in addition to sharing spectrum assets between operators who have primary ownership of these assets, a new regime for shared spectrum access is developing with a focus on bands occupied by government and military users, such as 3.5 GHz in the United States and 2.3 GHz in Europe. Dynamic spectrum access will allow operators to access spectrum on a secondary basis particularly for small cells that are used to augment capacity on targeted basis. While the regulatory regime for spectrum access is still under discussion, there is great determination to realize this approach by regulators who are eager to kick-start a new wave of innovation and its accompanying economic benefits.

Resources Sharing and the Emergence of the Cloud-Centric Virtual Operator

Developments in wireless network architecture towards virtualization and increased resources sharing, such as Cloud RAN where the radio access network is transformed into a common hardware infrastructure, would not occur in vacuum and can be accompanied by equally innovative business models for the MNOs and the (over the top) OTTs running on top. Given the substantial holdings of spectrum by a number of wireless players around the world, including spectrum that would be optimal for re-farming from alternative technologies (e.g. WiMAX, CDMA, etc.), there is a case to be made for the emergence of utility-oriented mobile Internet providers. Google, Apple, Amazon, Ali Baba, Tencent and various large-scale Internet and cloud players, have an opportunity to operate a virtually isolated network, or network service provider (NSP), within various disruptive business models ranging from device or applications priced-in bandwidth to

select volume unit billing models. These various ways of sharing spectrum and network assets with Internet players will form a new breed of mobile virtual network operators (MVNO).

In fact, and as the mobile Internet becomes an elementary expectation and as participation in the global conversation becomes more critical to the individual, the wireless operator market will likely evolve towards this position. This would come in as a handy deployment model due to the fact that the incumbent service providers cannot achieve the low cost of capacity required to enable this model.

Unlike the “cellular Internet,” the opportunity exists to develop a mobile Internet utility ecosystem that builds upon intelligent sharing of spectrum and network assets. It will enable business models that would drive revenue for the Internet players using both subscriber conversions to an ad-free service, and premium fine-grained advertising utilizing location, declared interests, and preferences. This revenue will allow bandwidth pricing of the service and allow for various models of revenue sharing with spectrum and network resources players. This will take advantage of the mobile Internet utility model to deliver access to Digital Divide or poverty-unconnected users with low cost devices and pricing models, made possible by the optimal spectrum and network resources sharing, as well as optimal arbitrage between available network and cloud computing resources. Finally, it will deliver computationally intensive cloud applications to the handset without consuming precious resources by taking into account the scaling characteristics of cloud computational models.

Take aways

The increase in bandwidth requirements of wireless services has paradoxically increased dependence on fixed-access infrastructure (fiber optical networks), and heightened attention on alternative complementary access schemes (Wi-Fi). This, in addition to developments in base station and mobile network architecture have led to the emergence of new trends in active network resource sharing that are complementary to ones we have witnessed over the past decade. Although various active resource-sharing models are possible, we anticipate that they will be mainly complementary to and built on top of the existing fiber/backhaul, Wi-Fi and passive DAS models which would be extended into active DAS, small cells and Cloud RAN architectures. Furthermore, as a direct continuation of this evolution, new spectrum and airwaves resource sharing will emerge. A direct consequence of this, in our opinion, is the rise of Internet and cloud-centric virtual network operators who will take advantage of optimized network sharing and wholesale delivery models, and introduce novel business, pricing and revenue share models that will constitute a significant disruption in how mobile Internet services are provided.

Acronyms

COTS	Commercial off the Shelf
DAS	Distributed Antenna System
FTTx	Fiber to the x
HetNet	Heterogeneous Network
LTE	Long Term Evolution
MNO	Mobile Network Operator
MVNO	Mobile Virtual Network Operator
NSP	Network Service Provider
RAN	Radio Access Network
SEA	Southeast Asia

XONNA PARTNERS

The Path to 5G Mobile Networks Gradually Getting There

Dr. Riad Hartani, Frank Rayal,
Mano Vafai, Richard Jeffares,
Mats Vilander, Dr. Dean Sirovica

January 2014

Setting the Scene

The years between 2005 and 2010 were perhaps more unique and interesting in the space of wide area wireless communications than any before. During those years, LTE was born as a fourth generation technology in as much of an effort to stave off competition from WiMAX as it was to provide a roadmap for operators who were in a quandary on how to recoup their sunk investments in 3G networks, not to mention the search for a killer app for data services. The introduction of the iPhone in 2007 and the advent of the smartphone created a sustainable and insatiable demand for wireless capacity which propelled LTE to be deployed on large scale by many operators around the world, creating for the first time a worldwide standard for mobile communications.

With data services here to stay, the question is then how can the mobile industry meet the long-term demand of subscribers? What are we to expect over the next 5 to 10 years, as far as mobile network evolution? Would we still witness a linear evolution from 4G to 5G, mostly lead by 3GPP/ITU-T specifications, or is it likely to be heavily influenced by the fast moving IEEE Wi-Fi standards evolution? Or even going one step further to anticipate mobile overlay applications to dictate how 5G gets defined?

We aim in this paper at addressing some of these aspects and flush out some of the fundamental architectural developments and mobile technology deployment models that one shall anticipate as we get into the era of beyond 4G and into the still yet to be defined 5G era.

A Historical Perspective

First generation wireless networks deployed in the 1980's were based on analog modulation. These include AMPS, TACS and NMT. In the early 1990's digital modulation was first introduced by GSM which became a de facto world standard and CDMA IS-95 (commercially known as cdmaOne) which took hold in North America, Korea and a few other markets. Third generation networks were trialed in 2000 and featured packet data service while voice service remained circuit switched. 3G is based on wideband CDMA which uses direct sequence spread spectrum techniques over a bandwidth of 5 MHz (effective bandwidth is 3.84 MHz) as opposed to cdmaOne which uses 1.25 MHz channels. 4G LTE systems are currently rolling out worldwide and feature a complete packet-switched function where even voice is packetized. LTE uses OFDM physical layer with a scalable channel bandwidth up to 20 MHz to deliver mobile broadband quality of service.

Definition of Generations

Is LTE a 4G technology? This question raises a question on how technologies are classified. If the benchmark for defining a generation is the set of requirements specified in IMT-Advanced by the ITU-R, then LTE in its early incarnation (i.e. 3GPP Release 8 and 9) falls short. However, if one considers the evolution of the architecture across the entire communication protocol stack, then LTE can be considered a fourth generation technology. Pragmatically, classification is based upon a set of rules, which brings different set of perspectives. LTE defines a new physical layer and flat-IP core network architecture and, from that perspective, is a unique generation fully distinct from 3G and earlier generations. The LTE roadmap allows it to meet IMT-Advanced specifications.

Evolution within a Generation

Every generation is born with a roadmap to improve performance over time. GSM was first designed to provide circuit switched voice and later incorporated circuit switched data service called GPRS, which promised to deliver peak rate up to 114 kbps. 3G first started with the promise to deliver 384 kbps downlink rate in Release 99, which was set to the requirements of IMT-2000. The technology evolved with successive generations to provide peak 42 Mbps in its HSPA Release 8 multicarrier version deployed by some operators worldwide (even higher rates are claimed by later releases, but the prospect of commercially deploying such releases is small as operators shift investment from 3G to 4G networks). Baseline LTE performance is that of 3GPP Release 8 which defines the first LTE release. 3GPP, the body responsible for standardizing LTE, has defined a rich roadmap of features to improve the performance of LTE to meet the targets defined by IMT starting with 3GPP Release 10 (commonly known as LTE-Advanced). On the core side of the network, a parallel evolution path has been taking its course, with an initial architecture based on circuit switching technologies in 2G, then hybrid circuit / packet switching technologies in 3G and a goal of a flat IP based packet switching technologies in 4G.

As of the time of writing this whitepaper, work is ongoing to define Release 12 and exploring the features required for Release 13. Yet, today, most operators around the world are operating Release 8 and 9 networks with certain features of Release 10 in service by a few network operators (primarily carrier aggregation to increase throughput).

Defining The Future Challenge

Every generation of wireless technologies is more spectrally efficient than the previous generation. However, the incremental improvement in performance between generations is shrinking. With every generation we get closer to the theoretical limit defined by Claude Shannon in his famous equation linking capacity, bandwidth and noise level. We are close to this limit with LTE which incorporates the latest of techniques designed to improve performance such as OFDM physical layer, convolutional turbo codes (CTC), multiple input multiple output antenna systems (MIMO), hybrid automatic repeat request (HARQ), and adaptive modulation.

Successive generations leverage wider channel bandwidth to deliver higher data rates. GSM used 200 kHz and 3G used allocations of 5 MHz. LTE uses a scalable channel bandwidth up to 20 MHz. However, access spectrum is limited, especially that in sub 3 GHz used by the overwhelming majority of wireless networks.

The twin challenge of limited spectrum resources and tapering improvements in physical layer capacity will define the next phase of developments in wide area wireless networks. With this perspective, 5G wireless networks are expected to be better defined and characterized by techniques that allow different nodes to coexist and collaborate among each other constructively to limit the effects of interference. 5G would also be characterized by incorporating spectrum in higher frequency bands. While the physical layer changed significantly in migrating between 1G through 4G networks, it is expected to take a less prominent role in 5G where it would still be based on some form of multi-carrier access scheme (be it OFDM or more efficient techniques). It is also expected that the core network would remain based on IP (Mobile IP specifically).

What is 5G?

5G as it stands today is not a defined technology or even a set of requirements. It is a reference in industry circles of what is beyond LTE that often refers to beyond 2020 timeframe (estimated deployment in 2020-2025 timeframe). Because 5G is in the process of being defined, there are many definitions and views on what 5G is and is not. What is certain is that the incremental improvements in the capacity of physical layer would not alone meet the demand for data services, nor would additional spectrum grants, especially in prime spectrum for mobility services (sub 2 GHz). Different techniques are required to improve the efficiency and capacity of wireless networks to meet future service requirements. 5G will focus on providing this gain through a number of features and concepts that have been around in research circles but have not yet seen their way to full commercialization. In fact, some of these features have actually have been defined in the standards, but 5G will take these concepts to a new level as the standard will be designed from the start to incorporate such features. Here we note that the challenge is often in implementing such techniques – standards do not define how a feature should be implemented. Increasingly in modern communication systems, implementation necessitates logic, which is defined in software. From this perspective, 5G will comprise a heavy element of software, both on the radio and core side of the network, that will differentiate vendor's solutions. 5G is therefore about the intelligent network where coordination and coexistence are the hallmarks defining the network of the future. This could potentially provide a great strategic advantage to leading equipment vendor and will in turn increase switching costs for operators.

5G Activities

The EU recently funded a research program under the name of METIS with a €50 million grant to develop 5G technologies and regain some of Europe's lost leadership in mobile communications. Some of METIS objectives include :

- 1000 times higher mobile data volume per area: network operators will serve many more users at the same time.
- 10 times to 100 times higher number of connected devices: new smart technologies will be invented to connect cars, appliances, and home energy and water controls.
- 10 times to 100 times higher typical user data rate.
- 10 times longer battery life for low power machine-to-machine communications: provide more autonomy on the move and lower energy consumption.
- 5 times reduced end-to-end latency for smoother interaction with bandwidth-hungry applications and less waiting time.

This is one example of what 5G can look like – but we stress that it is not a universal view. Other entities including vendors, operators and industry forums have their views. 5G is still too early a topic for standardization, but there are trends to follow in mobile communications that can give us a glimpse of the future. So, what can we expect to see in 5G?

5G Air Interface Highlights

Network Densification: Increasing the capacity of wireless networks by multiple folds to meet demand necessitates deploying cells with small coverage radius. This is likely to be achieved using different types of small cells. While the term ‘small cell’ often refers to a compact base station, it is used in this context to refer to any transceiver covering a small area. This transceiver can be a remote radio head connected through high-speed fronthaul system to virtualized pool of baseband resources, which is known as Cloud RAN (CRAN). The small cells can operate in different technologies (today Wi-Fi is prevalent as are 3G femto cells). Small cells can operate higher in the frequency spectrum to provide greater throughput.

Network MIMO: Coordinated transmissions from multiple base stations, or network MIMO, has been defined in LTE Release 11 as Coordinated Multipoint but not yet implemented. It is expected that 5G will include coordination as network MIMO reduces interference. Coordinated transmission helps improve cell edge performance in particular but requires fast connectivity between the transceiver nodes.

Massive MIMO: Massive MIMO involves a very large array of antennas at the base station to serve a large number of users simultaneously. Massive MIMO can work with centralized or distributed antenna systems and can operate with some form of coordination. Some of the challenges include logistical issues of how to pack many antennas on a base station site. Massive MIMO may be deployed on small cells operating in higher frequency bands, which become a more manageable proposition from implementation perspective.

Cooperative Networking: The networks of the future are heterogeneous that comprise different nodes including macro, femto, pico cells, relays and Wi-Fi cells. In such an environment, multiple nodes can cooperate to serve a device. LTE defines certain cooperation techniques such as ICIC (Release 8 and 9) and eICIC (Release 10 and 11). 5G would incorporate more advanced forms of coordination between nodes and technologies. For example, in a step to expand on this concept, a device can serve other devices should it have a good communication channel. This is termed ‘client cooperation’ and is sometimes referred to as ‘multi-hop communications.’

Cognitive Radio: Cognitive radio is a concept centered on agility of selecting the operating frequency band, channel bandwidth, and physical layer according to the environment, traffic load and other parameters. Cognitive radio enables accessing the same spectrum resources efficiently by adaptively identifying unused spectrum and adapting the transmission scheme to the requirements of the technologies sharing the same spectrum. By definition, cognitive radio implies the ability to sense the channel in order to adapt its transmission, which has proven to be a challenging task. Advances in cognitive radio technology would allow certain implementations to be incorporated into the 5G standard to increase the efficiency of spectrum utilization.

PHY Improvements: OFDM is a robust scheme for communication in fading channel. However, it suffers from certain inefficiencies. In the frequency domain, it has a relatively high side-lobe level and slow roll off. In the time domain, the cyclic prefix in LTE accounts for about 6.5% in overhead. Additional forms of multicarrier access schemes are under study including Filter Bank Multicarrier (FBMC) technology, which is a form of tightly packed FDMA carriers that results in greater spectral efficiency than OFDM.

Super Wideband Spectrum: Trunking theory shows that a wide channel carries traffic more efficiently than multiple narrow channels of similar aggregate bandwidth. Hence, a 20 MHz channel would have higher capacity than 2x10 MHz channels. Throughput increases linearly with available spectrum. Furthermore, spectrum in higher bands is more abundant than in lower bands. Using super wideband spectrum is another way to achieve the high capacity targets for 5G networks. In high frequency bands, directional antennas based on beamforming technologies would provide directivity and gain to close the communication link.

5G Core Network Considerations

The design of the 4G core network (EPC), as defined in the 3GPP EPC/SAE specifications, lays out the basis for a flat IP-based architecture supporting LTE and its evolution to LTE-Advanced, as well as the interworking with 3G and other technologies such as CDMA and Wi-Fi. As such, the evolution of the EPC/SAE is not expected to fundamentally change the overall functional architecture in terms of elements and interfaces, but will definitely change its implementation, scale, performance and programmability requirements. This is driven by anticipated deployment models that include supporting large number of end points as required by M2M and IoT applications, providing greater control to end users, enabling a dynamic interaction with the OTTs, optimizing for vertical-specific MVNOs running over the wireless network (which can be based on industry vertical models, such as e-health or automotive, or branded-device MVNOs such as an evolution from the Amazon Kindle model into a large variety of cloud managed branded devices, as an evolution of the Google, Apple and other application delivery models), and supporting small cells and Wi-Fi networks as a service. All this mandates increased flexibility and programmability within the constraints of lower total cost of ownership (both capex & opex) deployment model.

With this in sight, the mobile core network, including the EPC and the various components it interacts with, will evolve, in some specific instances, towards an architecture leveraging virtual machines (VM) and hypervisors technologies that run on premise or in a cloud environment. This architecture lays the foundation for a transition from dedicated hardware systems to SDN models to control the virtual environments and various components of the architecture built with NFV concepts. The key objective is to create highly scalable networks with a lower capex and opex than existing networks while introducing new service delivery models, as required by the emergence of new business models for mobile operators. In fact, a lot of these considerations are being experimented with already, where most of the focus is on validating early stage software implementation, integration into the back-end environment, refining migration strategies and developing fully interoperable multi-vendor implementations. In many ways, the core of the mobile network will witness a lot of the developments that have first happened in the data center, as far as virtualization and cloud deployment models.

EPC/SAE Implementation in a 5G Environment

The various elements of the EPC and complementary elements, for example IMS and billing/charging elements as well as the Value Add Services network, will progressively migrate, when the right conditions are set, from dedicated hardware to virtual machines. Initially, and as long as the software is centralized on the VMs, there will be no real change in terms of functional requirements.

Later, some functions, which are driven by services and deployment requirements, will progressively be built over virtual environments that are distributed over multiple VMs and in some cases run in private or public cloud environments. Getting to this stage will require a re-architecture of the EPC network through reconfiguration and adapted messaging over various interfaces, which is likely to require some new or adapted standard specifications. The new architecture will need to address various functional blocks of EPC and the elements it interacts with on the northbound, as an example, the interaction with the IMS VoIP (e.g. MTAS / IM-SSF / SCIM / P-S/CSCF related functions for orchestration, HSS interaction with the services layer, etc.). The key focus will be on addressing the performance, security, interoperability and QoS impacts resulting from this transition. The winning architectures would be the ones allowing a smooth migration that minimizes the disruptive impact and lowers cost.

The Path for Core Network Evolution

To illustrate the evolution of the core network, let's look at a brief description of how the design and implementation of some specific EPC components is likely to evolve from its current state in the next few years:

The Mobility Management Entity (MME) provides the overall mobility management and session management functions in the LTE network. The MME functionality would be one of the first functions migrated to a central or distributed virtualized environment largely driven by a novel set of service delivery functionalities.

The Serving Gateway (S-GW) provides the mobility anchor point for a LTE mobile device to access data services. The PDN Gateway (P-GW) provides access to one or more Packet Data Networks. Data path performance requirements, as well as the integration of functions that were adjacent to the packet core into the packet core, such as video caching, video transcoding/trans-rating and various stateful security considerations, will require the S-GW and particular P-GW functions to run over dedicated high-performance hardware for time to come. However, specific deployment models, such as dedicated vertical MVNOs, end-user controlled networks, dedicated M2M and IoT overlays will lead to the emergence of S/P-GW implementations running in virtualized environments, in either private cloud environments if controlled by the mobile operator or private/public cloud environments if controlled by the MVNOs or end users.

The Policy and Charging Enforcement Function (PCEF) is a part of S/P-GW and it enforces Layer-4 to Layer-7 Policy and Charging Controls (PCC) provided by the PCRF. This enables service based routing, packet forwarding, traffic shaping and policing. The PCEF functionality will follow the same deployment logic as that of the P-GW as it is seen as continuity to the latter's various functionalities.

The Policy and Charging Rule Function (PCRF) provides Policy and Charging Control engines for a service provider to define network/application service policies and charging rules to a subscriber or a group of subscribers. The PCRF, as a network-wide controller, will progressively run over VMs in either private cloud environment when under the control of the mobile operator, or possibly in public cloud environments when providing control function to overlays, OTTs and MVNOs over network resources.

EPC back-end and underlying IT transformation

Aside from the evolution of the various EPC elements, the 5G core architecture is envisioned to be most strongly influenced by the way the data and IT architecture around the EPC are likely to evolve. This would include all aspects related to data aggregation off the core and services network, network data storage and warehousing, data querying, as well as third party applications that run over the data warehouse, such as business intelligence, and the various APIs that would expose these data to third party applications. The overall IT architecture will leverage a lot of the virtualization, cloud and big data architecture models.

Mobile operators will find themselves radically transforming their IT architecture to accommodate this transformation. Some of operators, having already initiated specific IT transformation architectures based on SOA models where various elements of the mobile core interact seamlessly with other elements over dedicated information and messaging brokers, with ESBs as examples, will find it easier to migrate to the new virtualized cloud and big data based IT architectures given the fundamental importance that seamless, flexible and scalable inter-element communication will have in such architectures.

Below describes some of the major trends and a set of possible software implementation of such functionalities over the next few years. These are provided as examples, noting that various other implementation techniques are also available.

Business Intelligence: The architecture components are designed to provide off the shelf analytical components to fit in with minimal integration work. Building the business intelligence platform leveraging specific big data implementation (Platfora implementation framework, as an example) is a good choice and provides a good mix of integration within a Hadoop ecosystem and easy to use frontend for data analysis. This allows for a flexible ability to provide support for heat maps, charts and drilldowns to publishers.

Data Storage and Warehousing: Various implementation frameworks will be introduced. Hbase as an illustrative example here, is very efficient in fast time range scanning, time range queries, data drilldown etc. in the face of read only data with low throughput write data. Additionally, HBase supports quick snapshots and is ideal data warehousing platform. Data cubes stored in HBase allow cube operations such as pivoting and drill down via HBase. HBase is a good data warehousing option in terms of cost/performance for report generation. In a similar way, Hive on Spark allows for in memory queries for analytics that provide near real time analysis of data. This component will address SLA requirements of the reporting solution without having to implement the existing reports but with added performance. Additionally, this provides better data import/export than, for example, MongoDB noSQL solution with better performance for lower cost.

Optimization of Data Architecture Availability and Reliability: Here again, Hadoop 2.0 and upcoming iterations of Hadoop, as an example, will form the basis of the availability and reliability architecture. It supports distributed Jobtracker and high availability to Datanode. This avoids single point of failure for the Hadoop deployment. HDFS replication itself lends to high availability of data on file system. Zookeeper should be implemented with multiple nodes for high availability of cluster. Data policies for archiving and snapshots through HBase will provide reliability and disaster recovery options for the cluster. Data ingest is one of the critical first steps in achieving data consistency for analysis. Flume allows predictable and efficient data ingestion into HDFS file

system providing visibility into failures and improving performance of data ingestion. Missing data can be detected with custom plugins to Flume pipeline. Depending on the requirements, it is possible to use Kafka in the pipeline for reliable delivery of data for preventing data loss.

In a similar way, Oozie provides event based workflow mechanism for launching jobs in the event of data ingest into HDFS or HCatalog. Additionally, Oozie provides an easy way of specifying job workflow including Pig and Hive jobs allowing SLA specification for workflow. This implementation will allow better quality of data for reliable reports and better performance on scheduled reports as well as ad-hoc queries.

Data Management Performance and Scalability: Cloud based deployments will form the basis of the scalability models for the IT and backend architectures. Here again, and as an example, performance and scalability improvements are achieved using Hive on Spark. Linear scalability and performance with scale can be achieved by using the Hadoop 2.0 architecture as defined in the next section. This cluster is designed to be single cluster to support data needs that can consolidate all or some, of its data centers. Processes and policies in place for data lifecycle management for archiving, retention, compression and replication will allow for efficient data management with low overhead costs.

Network Monitoring, Metrics, and Diagnostics: Dedicated platforms will provide a dashboard for comprehensive monitoring of the cluster using frameworks such as Ganglia monitoring system or existing monitoring system along with Job profile and analysis. This in turn provides predictability to job completion times based on job profiles that provides excellent diagnostic capability for job performance and predictability. Fine-grained estimations on cluster usage by user, job type and time of day will allow for better policies for cluster usage and planning. The diagnostic insights lead for high performing jobs, better data design and lower failures in the cluster.

Data API and Exposure to 3rd Party Applications: Data API is a virtualization layer that hides underlying platform details and provides REST or JDBC interfaces for external interaction. There will likely be an evolution towards the integration of on and off premise Data API solution providers, natively or as SaaS model. Some solutions allow simplifies data import export to noSQL databases. These solutions can be integrated to provide a consistent data view to external actors. Application development using the data interfaces that are decoupled with data storage structure will lead to lower cost of maintenance and better integration with partners.

Real-Time Analytics: Real-time data processing has to accommodate high velocity data stream and process data in near real time for alerts and analysis. Real time processing systems, using frameworks such as storm and Kafka will allow for horizontal scaling, large-scale events processing, reliable data management and dynamic events handling.

Storm supports high throughput event processing and achieves reliability using Kafka for incoming data. Processed events can generate events that can be acted upon for real time processing by additional jobs. The processed data is persisted using HBase for efficient storage and can be combined with historical data in the cluster for generating reports at regular intervals. This real-time processing infrastructure will support mobile reports that are expected to be generated in near real-time, which in itself is a great value add as far as intrinsic business value.

Getting to 5G

The road to 5G begins with defining the requirements and objectives for the technology. Ongoing research and development helps define what technologies will be considered for inclusion in the future standard. Then standard activities will start to work out the details and achieve consensus among industry players. Different types of tests and trials will follow before commercial deployments. All throughout this time, the LTE roadmap will continue to evolve to include some new features that represent the precursor to those in 5G. For example, LTE Release 10 includes carrier aggregation which today scales up to 2x20 MHz for a total of 40 MHz of spectrum. Release 10 also includes eICIC techniques targeted enabling HetNet deployments in addition to many SON features that are required to enable operators deal with the complexity of large networks. Coordinated multipoint is defined in Release 11 but there has been no firm commitment for its deployment to date by any operator. On the core, backend and underlying IT infrastructure, a gradual move towards virtualization, specific functionality enablement in private/hybrid/public cloud environment, and in particular integration of big data analysis frameworks for the overall network data management, will start appearing in mobile networks core and services network environments.

What is clear in our reflection, is that we are at an inflection point in the mobile network and application development, taking advantage of fundamental technology shifts, but more importantly forcing new business and service models to emerge. In the years between now and when 5G becomes within reach, the LTE network will evolve to include many features that have been defined to date but not yet implemented, and would enable a new wave of mobile services that are yet to be envisioned. In all, it makes for a very interesting period as the next wave of innovation can raise the fortunes of vendors and operators who lagged and missed the LTE cycle and provide them with a new opportunity to displace today's leaders, while at the same time, creating new challenges to existing vendors and operators who have to face the threat of potentially disruptive technologies that would give a chance to their competitors to pull ahead.

Acronyms

2G	Second generation
3G	Third generation
4G	Fourth generation
5G	Fifth generation
AMPS	Advanced Mobile Phone System
API	Application Program Interface
CDMA	Code Division Multiple Access
CoMP	Coordinated Multipoint
CRAN	Cloud RAN
CSCF	Call Session Control Function
CTC	Convolutional Turbo Codes
eICIC	Enhanced Inter-cell Interference Coordination
EPC	Enhanced Packet Core
EFB	Enterprise Services Bus
FDMC	Filter Bank Multicarrier
FDMA	Frequency Division Multiple Access
GPRS	Global Packet Radio Service
GSM	Global System for Mobile Communications
HARQ	Hybrid Automatic Repeat Request
Hadoop	Open Source Software Framework – Hbase, Hive, Zoo-keeper constitute some of the projects within or related to the Hadoop framework
HDFS	Hadoop Distributed File System
HetNet	Heterogeneous Networks
ICIC	Inter-cell Interference Coordination
IMS	IP Multimedia Subsystems
IM-SSF	IP Multimedia Services Switching System
IMT	International Mobile Telecommunications
IoT	Internet of Things
IP	Internet Protocol
IS	Industry Standard
ITU	International Telecommunication Union
JDBC	Java Database Connectivity
LTU	Long Term Evolution
M2M	Machine to machine
METIA	Mobile and wireless communications Enablers for Twenty-twenty (2020) Information Society
MIMO	Multiple Input Multiple Output
MME	Mobility Management Entity
MTAS	Multimedia Telephony Messaging Server

MVNO	Mobile Virtual Network Operator
NFV	Network Function Virtualization
NMT	Nordic Mobile Telephone
noSQL	Not Only Search and Query Language
OFDM	Orthogonal Frequency Division Multiplexing
OTT	Over The Top
PCC	Policy and Charging Controls
PCEF	Policy and Charging Enforcement Function
PCRF	Policy and Charging Rule Function
PDN	Packet Data Network
P-GW	Packet Data Network Gateway
PHY	Physical Layer
P-S CSCF	Proxy / Serving Call Session Control Function
RAN	Radio Access Network
REST	Representational State Transfer
SAE	System Architecture Evolution
SCIM	Service Capability Interaction Manager
SDN	Software Defined Networks
S-GW	Serving Gateway
SLA	Service Level Agreement
SOA	Service Oriented Architecture
SON	Self Organizing Network
TACS	Total Access Communications System
VM	Virtual Machine
VoIP	Voice over Internet Protocol

XONNA PARTNERS

Data Sciences Focus Mobile Eco-System Contributions

Dr. Riad Hartani, Dr. Alex Popescul,
Dr. James Shanahan

November 2013

Data Science Evolution – A Brief Revisit

As a team, we have first contributed to the field of artificial intelligence and overall data sciences since the early 90s, with early work in the area of machine learning, multivalued logic and neural networks. We then followed the evolution of these data analysis and knowledge discovery techniques over the years, as algorithms became more elaborate, computing models more efficient, and live data generated and collected at increasingly higher rates, often for completely novel applications.

Very recently, our focus has been on analyzing applications of recent techniques such as deep learning and the various additions to random forest and gradient boosted decision trees to practical industry problems. Over this 20 years lapse of time, one thing became clear to us: Data Science has made the big leap of being a research area for a select few applications, to a set of tools, accessible in various shapes and forms to various industry verticals, and optimized to resolve some of their more challenging problems. In this paper, we synthesize our experience on the experimental front through a recent case study, applying and customizing select advanced Data Science algorithms to a new set of Internet services applications.

Specifically, we focus our interest on real world case studies in the realm of mobile and cloud networks optimization, and corresponding business intelligence models running on top of these networks. In fact, up to today, commercially available data analytics products on the market have had the following shortcomings: (1) a limited scalability of for the data collection models as measured in terms of data generation, collection and storage (2) a lack of efficient machine learning and predictive modeling algorithms to process collected data in real time, (3) an open loop data analysis feedback, that is not dynamically correlated with the operator's business logic and (4) computing and pricing models that are based on centralized localized processing in operators' IT infrastructure, that is not taking advantage of the pay- as- you- go cloud- based computing.

Our stated goal over the past couple of years was to address the above shortcomings, taking advantage of technology innovations recently introduced, and just now getting to a sufficient level of maturity to be commercially applicable to mobile networks data analytics, taking into account operators' business logic goals in mind.

Data Science in a Mobile World – Why now?

To achieve the stated goals above, we opportunistically leveraged the fact that three fast converging concurrent trends. A brief snapshot is presented here.

First is the maturity of Data Management models.

We are witnessing the fast adoption of novel architecture to store and access large data sets (Hadoop, MapReduce, HDFS, Yarn, etc. – commonly known as Big Data models), as well the commercial availability of various cloud deployment architectures (OpenStack, vCloud, Cloudstack, AWS, etc.). This is removing significant logistical obstacles to embracing management of large data structures. The move is likely to be even more significant moving forward, given the immense number of contributions of the open source community in this area (as an example, we were part of a 3000 contingent at the last Openstack summit in Hong Kong – November 2013 -, the largest

ever, which shows the strong interest of the computing community). Key here is convergence onto universally adopted platforms versus what was before seen as a proliferation of diverse platforms.

Second is the evolution of Data Sciences.

This applies to the large set of data analysis models in a broad sense, and specifically machine learning and mining algorithms that are more accurate and computationally tractable, leveraging distributed cloud-based computing models. Current developments in Deep Learning, for example, illustrate well how an older field of neural networks achieved breakthroughs in accuracy when its algorithm improvements were fueled by much increased computational power. Taking advantage of the introduction of new computing models, such as algorithms parallelization, GPUs and alike, then porting that to distributed cloud compute models, not only the existing algorithms have been optimized to run better and faster, but a number of additions and optimization have been developed and run in a computationally tractable way.

Third is Data Availability.

Leveraging software and hardware architectures that are increasingly scalable to selectively and dynamically process large volumes of data, relying on various models of data capture, via sensors, devices, and management modules. Larger data sets influence algorithm choices by easing the risks of over-fitting, which leads to better generalizable insights. The sheer size of data available is likely to increase, either as front-end data in real time or backend data stored as historical patterns. In the networking world specifically, hardware-based data-path architectures have evolved in a way that allows for data to be captured fast enough for deeper analysis, and software-based management architectures in a way that data can be queried, received and presented to relevant data processing models.

It is in fact, the first time ever, that such trends are coming together, which opens up the opportunity to leverage the vast amount of real time users and services data available for processing, through a correlation of to its underlying business processes, to optimize bottom line business logic, and dynamically derive new revenue streams and optimize existing modes of operations.

One specific real-life case study we have recently worked on with Tier 1 global mobile operators is briefly explored. It sits in the context of optimizing and monetizing their mobile data along select dimensions. Various similar case studies, in the areas of mobile data fraud and revenue assurance, public cloud migration enablement with underlying performance measurement and enforcement, as well mobile payment models optimization have been or are being worked on. This builds on very similar set of tools developed by the team over the last few years, in the world of digital and online advertising, web search optimization and related topics.

A Mobile Network Optimization Case Study

The Data Science solution we have worked on is inherently modular, and part of a more elaborate solution umbrella, composed of:

(1) A hybrid local/cloud based data gathering and storage, leveraging novel techniques optimized for the variety of data models. Adaptations of Hadoop-like models and their underlying MapReduce computing paradigm for large scale distributed file systems, are leveraged to present the various

data sets, that normally gathered in silos into a common data representation accessible to data processing models and

(2) A set of machine learning and data mining algorithms, specifically focused on clustering and predictive modeling in high dimensional spaces based on imprecise, uncertain and incomplete information, efficient statistical data summarization and features extraction algorithms as well as large scale real time data streams management. These tools will be at the core of the processing engine, and will aim at deriving optimization to the existing business logic and augment it with new revenue generating business logic, which would be mapped to a set of new revenue generating services.

3G and 4G networks are built over flat IP packet based networks. With the flexibility and scalability of IP based networks and services, comes the requirements for more stringent traffic and resources management mechanisms, and underlying challenges, unseen in previous circuit based switching technologies (for both user data where TDM circuits are replaced by IP / MPLS sessions) and control data (where SS7 is progressively replaced by SIP and Diameter IP based signaling).

The new architectures introduce various network elements in order to tackle such challenges. This would include data path processing models such as Deep Packet Inspection devices, used for marking and rate limiting traffic, to data compression/rating devices used for video optimization to topology and state aware control plane devices such as PCRF engines and SON resources load balancing engines among others.

In order to optimize customer user experiences (Quality of Experience (QoE): defined along various KPI metrics as perceived by the user), 3G/4G networks require the introduction of more sophisticated predictive, preventive and/or corrective resources management models in the networks. This is specifically where we have introduced novel data processing models, leveraging machine-learning algorithms, and demonstrated their value. As such, a real world traffic control scenario is developed, addressing a very specific problem that is causing major challenges in mobile networks today. The problem is formulated as follows: How to maximize the aggregated users QoE utility function over time, based on observation of real time and batch historical network level data measurements, and enacting semi real traffic control mechanisms in specific network enforcement points (either directly through dynamic provisioning or via a policy proxy function, such as a PCRF spell out for non-specialists).

This problem is instantiated via the following specific case study: QoE of the users is, in this case, modeled as the proportion of users traffic facing admission control rejection on setup (during the signaling phase between the mobile user device and the Radio Network Controller in 3G networks or equivalent in 4G networks between the user and the mobile network packet core), the network measurements are observed off the radio base stations either directly or through some level of aggregation, and the enforcement policies are based on pushing rate limiting decisions on the data-path, as well as other mechanisms focused on the video angle for transcode/transrate, etc..

The following assumptions have been made, for illustration purposes, but without impact when a more elaborate network model is taken into account (example: various dimensions are observed

to determine congestion levels in different parts of the networks and different multi layer mechanisms are enacted to push policy enforcement decisions at various interfaces): data off the radio base stations are captured in various network conditions, and in various network locations, where no resources management model is enforced (besides the ones intrinsic to the radio access and core access intrinsic resources management models as used in a standard configuration). Data is modeled along an input / output dimension space, where the input shows a multidimensional aggregate packet setups entering the network over time, and the output showing the blockage levels over time. Machine learning algorithms (time series, neural networks, deep learning models explored) are trained on such data to model this function and provide an approximation of such function over time. The learning model would optimize the time horizons in the past over which the data is read (as input to the approximation function) and the time horizon in the future over which the function is being approximated.

Thresholds are defined where the network entry acceptance reaches some configured level would be identified as a threshold over which resources management mechanisms would need to be pushed down the packet core to reduce traffic volumes (either per user or an aggregate across users). The higher the projection of blockage levels into the future, the more aggressive the rate limiting would have to be.

Rate limiting functions would force Internet and private traffic (assumed to be a mix of TCP and UDP, with different rules applied to each) levels to drop by a well-defined function over some defined time interval (delayed in time versus the time where the policy action is pushed down). Assumptions are that this decrease will result in a step function reduction for a set of recommended policy actions. It is also assumed that such aggregate reductions in traffic would cause a slowdown in non-interactive traffic without affecting the interactive traffic, and hence marginally affecting the QoE utility function.

Based on the assumptions above, the machine learning models, coupled with the closed loop control feedback model would demonstrate the following:

The traffic projection is modeled with a sufficient accuracy, over time, leading to an appropriate approximation of entry into the network rejection levels. The model would run based on input data, and as soon as the projection shows a high level of rejection in the future time horizon, a control policy action is pushed onto the network. This control policy action would then force traffic levels down, and as such feed updated data into the prediction model. The overall system would run with this closed loop feedback and overall maximization of the utility function is proven, while the overall network stays in stable conditions.

Conclusions and Key Take-Aways

A brief description of some of the Data Science applications to mobile networks have been highlighted, as a way to demonstrate applicability and value of such techniques in the real world. Specifically, one demonstrates that existing vertical industries (mobile telecom world in this specific case), that have historically been fairly slow moving in terms of pushing new data analysis techniques, are starting to get disrupted. Disruption in this case is beneficial, as it will likely converge on making operations way more efficient, build a platform for new revenue generating services and push towards a new generation of players, taking full advantage of the potential of Data Science models.

Xona Partners team, with its diverse technology expertise in the Data Science space as well as select industry verticals, along with its global insight into new business models developing across the globe, has been working with select players, in a win-win model, to solve some of the leading multinational pain points – or allow them to develop an edge in what is, and will increasingly become, a highly competitive play, where winners take it all.

XONNA PARTNERS

The Internet Eco-system Value Chain:

It's Always Greener On The Other Side.
Or Is It, Really?

Dr. Riad Hartani

October 2013

Over the last few months, I have had the pleasure to chair and moderate various global internet, information and communication technologies events, including 4G Mobile, Submarine Networks, Wholesale Operators & MVNO evolution, Technology Investment, Startups Innovation, and Telecom Services Strategies conferences and workshops around the world. This brought up, as it always does, insightful debates on the state of the affairs, market disruptions, underlying challenges, and provided a glimpse of the road ahead.

Recently, following a pause and reflection, some valuable take-aways surfaced clearly, and can be summarized as follows: the various players in the communication and Internet technologies value chain are no longer working synergistically as complementary players, as we have always assumed. In fact, they seem to be working in a “your pain is my gain” mode, as if the whole thing was a zero sum game scenario. What makes things look even more worrisome is that each layer in the value chain, seems to look at adjacent layers with envy, as if things were better and greener on the other side.

A brief synthesis of some of the most argumentative debates I have been part of, provides an overview of why this seems to be the case. Specifically, here is a re-run of four key panels among players in the same peer group debating the state of their industry: over the top (OTT) players, fixed and mobile telecom operators, infrastructure equipment vendors, and venture and private equity investors in telecom infrastructure and underlying technologies. In an ideal world, investors would be funding telecom operators who in turn, as customers, generate revenue to equipment vendors and serve as an infrastructure platform to a variety of OTT players with end users such as consumers or enterprises sitting at the top of the value chain pyramid. This would have them all share the risks and rewards of a healthy and growing global Internet and communication sector. But these were not necessarily the conclusions reached by these four groups of players, during the various interactions. Let’s look at the market landscape from their point of view.

First, the network operators.

This includes both the fixed and mobile access operators. Interestingly, they are in strong agreement on the threat of reduced ARPUs and declining profit margins. They point the finger at OTT players who are grabbing most of the value of high margin services (Skype/Viber in voice, whatsapp, wechat in SMS and Google in advertising), an unfriendly regulatory environment (cable operators in the US with network fairness requirements, mobile operators in Europe with the upcoming roaming regulations), as well as disruptive and rapidly standardized technologies that didn’t provide them enough time to capitalize on large sunk investments (fast track 2.5G to 3G to LTE migration). Such factors are compounded by a number of other factors which include macro developments that have made private equity investments in network infrastructure a rarity (very few greenfield operators are funded this way at the moment), geo-politics that lead to a risky “trans-border acquisitions as an expansion” model, and a single-vendor strategy by equipment vendors that is deliberately designed to lock operators into solutions from the large vendor/system integrators (as a result of vendor financing and managed services offering by the large network equipment vendors), which consequently make solutions roadmaps too rigid, products in the long run too expensive to deploy, stifles competition and innovation, and means that execution is slow, costly and constrained. In summary, network operators view their pain as the OTTs’ gain; a pain that is caused by equipment vendor strategies and a lack of investor appetite to support aggressive business models.

Second are the OTTs.

They come in various shapes and forms, but overall they build their service on top of fixed and mobile networks and address the same end-customers. In the view of OTTs, the blame is on the winner takes all model (a tiny percentage of OTT application providers make it to market successfully). For the vast majority of OTTs who don't end up as ultimate winners, and even for some of those who were able to scale and win the game, the blame is on what is called under the top players (UTT), or in other words, the network operators. The operators are blamed for being too slow to adopt partnering models (as in the case of mobile payment and mobile advertising) that allow OTTs to better leverage their assets via Open network Application Programming Interfaces and dynamic interaction models (slow adoption of efficient automated services provisioning models, as in cloud orchestrators or network layer Software Defined Networks). They are also blamed for shying away from information sharing models where network intelligence is provided for OTT differentiation (Business Intelligence contextual data leverage), and for pushing back at developing customer leverage models that would make it easier to develop new user interaction models for various verticals such as mobile payment, m-health, and automotive MVNOs as current examples.

OTTs, as large scale software application providers over public/hybrid cloud infrastructures, view the equipment vendors' slow motion towards network provisioning models that leverage agile virtualized networking architectures, as a serious challenge, going against the OTT ultimate dynamic user-controlled network resources. As such, the challenge for the OTTs in growing their business is basically to figure out how to leverage – if not exploit - the UTTs, and to some extent, leverage the network vendors' infrastructure products beneath it. The more successful they are at this, the more they divert resources from the less than adequate deployment models that are already causing the fall off in Telecom Operators' ability to partner or create innovative services.

Third are the infrastructure equipment vendors.

They can be classified in two categories: the large vendors, who provide turnkey solutions and the smaller ones, who provide niche plays in early deployment cycles. Consensus is prevalent here too, albeit perceived differently by each group. Vendors see operators' slow evolution and technology adoption models as the main reason for their revenue and margin challenges (slow migration to IMS/RCS based services, Network Function Virtualization deployment models, large scale M2M, etc.). For example, vendors in the LTE ecosystem have to go through the operators' 3 to 4 years cycle in adopting new technologies such as small cells, advanced backhaul and core network architectures, as well as slow integration of OSS into the overall IT enterprise architecture. At the same time, vendors in the application eco-system space face a push back against new overlay payment, advertising, and M2M deployment models. Vendors point to the lack of operators' fast adoption of technology and business models as the main reason for having venture capital and private equity funds shy away from network infrastructure investments which in turn slows down innovation and makes differentiation inherently difficult. Moreover, large vendors would add to this the fact that mobile operators, via their slow moving decisions, have all the financial leverage to play vendors against each other and, as such, significantly affect their margins and business models. The equipment vendors' CFOs and CEOs point out that network operators are not providing them with the revenues they deserve and that investors are not supporting the needs of their long sales cycle.

Fourth and last are the investors in network infrastructure and technologies.

Here again the perspective is fairly consistent. Returns haven't been what they should have been over the last decade (very few greenfield operators IPOs, rare success of trans-regional acquisitions). As such, the appetite for risk and new investment is fairly low at best. Large investors in particular see operators as investment vehicles who haven't managed to turn their cash generating models into high growth engines (utility models orientation), who have been shying away from new business models that could have shared the OTT service revenue (rare voice and video OTTs partnerships), and who have had limited success at growing organically by leveraging their customer base, geographically through acquisitions, or strategically through moves into adjacent markets such as online advertising and high value vertical markets analytics. Investors view the equipment vendors' push for open source models (Openstack cloud, Open Source controlled networking hardware, Hadoop framework for Big Data Compute) as a play against forecasted returns of custom design players. Investors also view the large vendors' links to operators as an impediment to high returns on investments in niche advanced technologies. As a result, with a paramount interest in return on investments through successful exits, investors view the infrastructure play as an industry in need of a different breed of network operators, as well as vendors in the infrastructure and OTT spaces needing more robust models for monetizing their product portfolios

What strikes me is that all these four groups were all adamant in agreeing to one thing: we are in challenging times mostly because the adjacent value chain-players are basically the embodiment of the exact threat we face, and our way out is to get a piece of the pie that these adjacency players are going after. In other words, our gain is their pain and our pain is their gain. What no one seems to be highlighting in the turmoil is that this pie is a growing pie, and that the rate of such growth is function of the synergy the various players can build among each others with the right business models. Of course, some of this logic is in the heads of the various players, and small start-ups tend to operate under this assumption most of the time, but still, the focus seems to be on laying blame as opposed to friendly – or synergistic - leverage.

Lets look at a couple of examples to illustrate how a friendly partnership can be leveraged. A strategic investor in telecom infrastructure technologies would see an increase in ROI when a mobile operator puts in place an architecture to monetize mobile data (as some operators are aggressively working on in the US right now), which would open new opportunities for solutions vendors and at the same time provide mobile OTT players with better returns on real-time online advertising models. A similar example would show how an investor into an optimized Hetnets (Small Cell / Wi-Fi) offload infrastructure (as some private equities are looking into in North America and Asia) would lead to a much better ROI for a neutral fixed line operator, which would in turn open up room for optimized network sharing vendor solutions, and through underlying novel business models, would open up the door for video OTTs to optimize their offering into a bundled wholesale and revenue sharing model. The growth of the Hetnet is an interesting example. It will clearly be an area of major investments, and leaving it to the operators means it will happen in an ad hoc way, driven by the trade off between fear of the competition getting there first and the constraints of their immediate budgets. But Hetnet expansion is an area where there is a clear win-win play for the joint community of investor, small cell vendor, infrastructure vendor, operator, and OTT service provider. The faster this gets rolled out, the more the consumer

will use it, the more revenue will be generated for all concerned. If there were some vehicle whereby each player could participate in accelerating the deployment, ubiquity and standards for the Hetnet in a very large sense of the term, then the inherent dynamics of the total ecosystem working synergistically would mean lower costs, and better and faster revenues and thus ROI across the board. The question then is why is this not happening, or not happening fast enough? Is it really now a zero sum game?

So, as a network operator, OTT, network vendor or Investor, one would ask: Is it greener on the other side? Do the network operators have it better than the OTT players? Do the OTT players have it better than the investors in network infrastructure? Are the equipment vendors hurting because network operators are squeezing their revenues? No certainties, but what is clear is that we are at some major industry inflection points, as far as business model change and underlying technology innovation is concerned.

For the industry player that embraces this change, we will witness a fast mover advantage winning big scenario: mobile operators aggressively moving into adjacent markets, fixed operators developing new Internet-centric and enhanced infrastructure sharing models, data center players scaling optimized cloud delivery models, video OTTs pursuing smart operators partnerships, vendors leveraging advanced integration of IT and network technologies.

The new models coming into place are predicated on making the pie bigger! One way or another this Internet, Telecom and Information Technology pie will get bigger, and fast. Sometimes desperate times call for desperate measures, and as such, the arrival of a fast track vehicle to embrace innovative business and technology models is what we are about to witness.

Key learning for financial advisors and investors.

So, what would be the bottom line for TMT investors and their advisors? Our belief is that they would need to shift focus from the limited appetite for in-market consolidation and trans-border acquisitions to portfolio rationalization and investment in synergetic adjacencies.

Private equity investors would need to capitalize on the emergence of such disruptive evolution models and thereby reduce the availability of investment in more classical growth models. Equity investors instead of polarizing between investing towards high growth in revenue or high ROI need to focus on companies that strike the right balance and sustain growth in revenue and ROI over the long term. It remains my firm belief that companies that are increasing the pie will be the ultimate winners!

XONNA PARTNERS

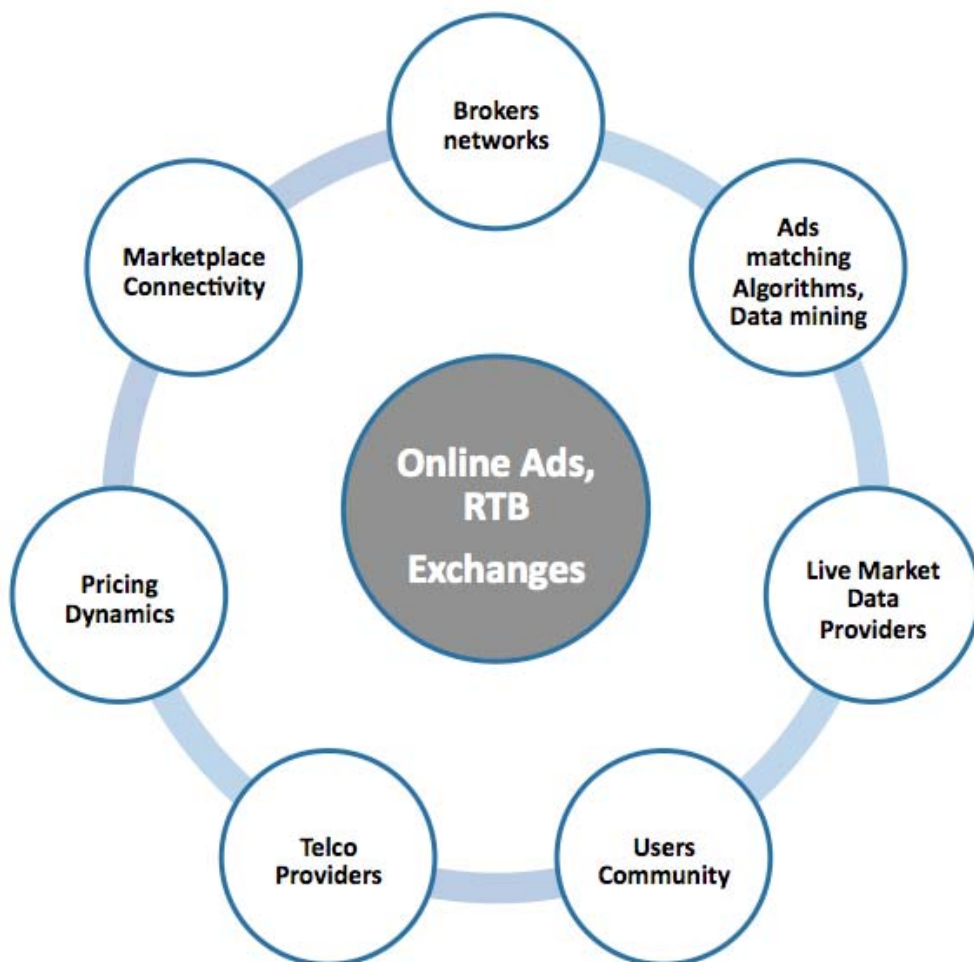
Online Advertising Real Time
Bidding & Opportunities for
Mobile Services Providers

Dr. Riad Hartani, Dr. James Shanahan

October 2013

As online advertising business continues its growth path (providing the core of the revenue streams of leading Internet players such as Google, Microsoft, Apple, Facebook, etc.), and mobile operators aim at sharing some of the revenue pie, along diverse business models, what would be the nascent opportunities for mobile service providers? We revisit the possible opportunities that online advertising evolution, and specifically the emergence of real time bidding (RTB) and growth of mobile advertising, opens up for mobile operators, in terms of business models and underlying solution development and deployment.

The online advertising eco-system is diverse and complex. Below is an illustration showing some of the various industry stakeholders. It shows the various components of the eco-system and the necessity of having all the dependencies taken care of to build an efficient mobile and real time advertising solution.



The ecosystem is fast changing and becoming increasingly competitive. As an example, the Ad exchange players, integrating DSPs (Demand Side Platforms) /SSPs (Sell Side Platforms) and the recent emergence of at least 2 major players (Facebook and Amazon) bringing to market their own solutions recently, in addition to the more established Ad Exchanges such as Google/DoubleClick, Yahoo/Right media and Microsoft/Appnexus.

Overall, the key problem to be tackled is to increase the revenue per user for advertisers, with mobile operators aiming at capturing a piece of such revenue, either through a revenue sharing model or through a direct revenue generation.

It's the revenue per user that would form the main metric of optimization, when comparing these numbers for various OTTs (Over the Top). Having mobile operators increase this revenue through various schemes is what would form the basis for new services, or new business models, with a direct consequence on the products and solutions strategies of the mobile service providers. This as well as optimizing bids of DSPs through RTB on Ad exchanges, with the complementary goal of optimizing CPMs (Cost per Mille), CPAs (Cost per Action) and CPCs (Cost per Click)

Mobile operators globally are in a phase where two options are put in front of them: either to optimize their networks to becoming a mobile broadband path, with no or little plans to share a piece of the revenues derived by the OTTs, or to position their network, selectively, within the overall OTT value chain, to share a piece of the revenue streams. This is also the case in the context of mobile advertising, where some operators, have been, and are still, working on defining their own approach to this market, now that mobile devices penetration is high, smartphones/tablets offer screens large enough for advertising and revenue streams off mobile advertising are seen as a good alternative to declining revenues in traditional voice services.

Few things play in the mobile operators' favor, including their existing relationships with customers, their existing SMS/MMS campaign based services offering and most importantly, the vast amount of customer data, that is of high value to advertisers. Few things however, remain challenging, including the fact that mobile advertising has never been in their DNA, hence requiring transformation, the fragmented nature of the customer base targeted by advertisers and the fact that OTTs (such as Google with the AdMob acquisition, apple with iAd products, etc.) have gone after this market very aggressively, making it difficult for new entrants to come in.

Multiple options are being considered in terms of how to approach the mobile advertising market. They are described below.

- Mobile networks directly acquiring mobile ad networks to build a direct presence in this space

This is the case of the largest mobile networks, with an aggressive push towards mobile advertising. The best example in this situation is Singtel through the acquisition of Amobee (mobile ad serving platform) and taking a significant stake into Nexage (a mobile Ad Exchange, with DSP/SSP integrated). In a similar fashion, both Telefonica and Vodafone have taken stakes into mobile Ad serving and Ad exchanges companies, which provides them with the option of building a business upon these technologies.

- Mobile networks partnering mobile ad networks to build a direct presence in this space

This is the alternative approach that some mobile operators have considered, as a strategy to approach the mobile advertising market. In some cases, this is a complement to going with the first option above as well. In this case, we can note the case of Telefonica, Vodafone, Etisalat (based on Alcatel Lucent solution), America Movil (based on myscreen solution), 3 Group (based on Rhythm solution) and Verizon Wireless. In these cases, the platforms are owned and managed by the partners, but through a well-defined partnership model with the mobile operators.

- Mobile operators building their own mobile ad platforms to compete with mobile advertising networks

This is the case where operators have gone into designing and implementing their own mobile ad solutions, such as Ad servers and to some extent Ad exchange, DSP/SSP platforms. This is still in early stages of development. Examples include AT&T and NTT Docomo. In some cases, operators have been working on sharing common co-developed platforms to address the fragmentation problem and increasing the size of the customer base and having it approach the addressable size, as seen per an OTT. This is specifically the case of small mobile operators who would need to join efforts to get to a sizable customer base, such as in HK, Singapore, Taiwan, UAE as examples. It is unclear as of today if such a strategy will be conclusive.

As a complement to such models, some operators are right now looking at having their own mobile ads integration within their applications app stores, as a way to counter initiatives such as the ones used by Apple via iAd and similar nascent alternatives by Google/Android and Microsoft.

- Mobile operators focused on defining new business models leveraging mobile advertising without directly managing the mobile advertising process of buying, selling and inserting ads.

This is the case of large number of mobile operators, and is in some cases done in conjunction with one of the 3 options above. In most cases, this is build upon the existing operations process of mobile operators, such as performing content re-formatting based on screen sizes/formats, augmenting their billing models to accommodate mobile advertising insertion models, augmenting their marketing campaigns with mobile advertising related information, leverage some of their data warehouses information to be exposed to the mobile advertising eco-system running on top of the network, and integrate mobile advertising with their content distribution networks (such as IPTV, etc.), gaming networks. Some new business models are emerging in this context. Examples include Blyk MVNO as used by Orange, KDDI's own ecosystem, etc.

It is worth noting that in these various models, mobile operators aim at inserting themselves into the mobile advertising value chain from different angles, based on a strategy that is optimal to them. One should note that although various models are being considered, various challenges exist. This would include the fact that mobile operators have little experience dealing with the various actors of the mobile advertising eco-system, don't have their customer data optimized for efficient targeting, are not used to working based on unspecified 3GPP standards, have to address various privacy considerations and are very careful getting into customer expectation management challenges that mobile advertising would cause.

Two key conclusions could be derived: first, the fact that the strategy to take in terms of development and business plan is not trivial and requires an in-depth review of this market, the pros and cons of each approach and more importantly a crisp mid to long term view of the customer landscape and approach to market. Second, the fact that no action plan and no decisions on this front, could quickly lead to striking out the chance of being a player in this market, with all the undesired consequences.

It is our belief that the network solution provider business, along with the mobile operator eco-system, will ultimately address the new opportunities offered by the evolution of mobile advertising and other OTT offered services over the next few years. A new landscape is likely to emerge, taking advantage of colliding large and complex eco-systems. Hence our most important recommendation to network solution providers is to take a systemic view at such evolution, be open to disruptions, manage risk return equations and converge on a clear strategic plan on how to address these opportunities.

XOONA PARTNERS

Technology Cycles Deja Vu

Dr. Riad Hartani,
Frank Rayal

September 2013

The recent string of technology IPOs in areas as varied as network security, social networking, mobile advertising and the likes, points to an interesting time in the technology innovation eco-system. It basically highlights the rapid go to market of technologies, that in essence, build on top of the latest disruptive business and technology that have reached critical mass, and are, as such the culmination of the actual technology innovation cycle. This is to last few years, but what this also says is that, the disruptive technology inflection point that would hit us prior to the end of this decade, is now scattered throughout various technology labs and startups in early stage of incubation.

This is by no means new, and in fact, its part of a cycle that has been reinventing itself at regular frequencies. A bit of retrospective in terms of technology development cycles, shows us that investments poured into different focus areas over time, leveraging the disruptions and building the basis for the next ones to come. Looking back, one can trace back to some major disruptive IPOs that uniquely highlight the stage of evolution along the technology roadmap, and pinpoint specific inflection points in the evolution cycle. First, we have the IPO by Microsoft (1986, operating system; software on personal computers) followed by Cisco (1990, switching equipment; infrastructure vendors), Netscape (1995, web browser; Internet as a media), Google (2004, search engine, core Internet), and Facebook (2012, social media; Internet as a platform). Every major IPO was preceded with a period hyper-activity with several companies vying for prominence. The market would be fragmented, divided between a varying numbers of companies and accompanied by a considerable level of speculation. The post-IPO landscape becomes more tame with a few leading companies controlling the lion's share of rewards.

Where we stand today in the cycle is a continuation of what started earlier this decade: optimization of the Internet as a platform. The recent M&A and IPO events are a good manifestation of this trend, which should continue for the next 2-3 years. This trend is being reinforced by the continued convergence towards everything-mobile, geography-independent computing, and everything-connects-to-everything phenomena, which place us in the phase of platform optimization and monetization of the converged fixed-mobile Internet. Some of the features of this era include:

1. Increased elasticity of the network through progression of virtualization and software defined networking through the compute, storage and networking chain to accommodate different user-controlled services
2. Transition from "send and receive" information model to a highly interactive model between users and content.
3. Form large-scale sharing platforms and social networking applications to create personalized, user-centric and controlled social ecosystems
4. The return of pragmatic artificial intelligence, through its manifestations of machine and deep learning, leveraging accessible data sets and tractable computing
5. Increased integration and programmability of silicon as a high-speed computing operating system underneath a layer of compute and cloud operating systems.

6. Growth in the 4G mobile Internet eco-system, with particular emphasis on the increased interaction between the network and over the top applications.

The above list can be boiled down to a few key characteristics that include: data, intelligence, information, management and optimization. In short, the focus will be ever more on creating semantics off data followed by business models that leverage these data semantics to monetize the converged Internet. As such, players that will embrace risk calculated technology and business model changes, we will witness a fast mover advantage winning big scenario: examples would include mobile operators aggressively moving into adjacent markets, fixed operators developing new Internet-centric and enhanced infrastructure sharing models, data center players scaling optimized cloud delivery models, video OTTs pursuing smart operators partnerships, and networking vendors leveraging advanced integration of IT and network technologies. This time, as it was the always the case before, those embracing change would be the leaders to stay, and others would be absorbed or disappear. Put simply, just like basic genetics!