XONA

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 are 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

Xona Partners (Xona) is a boutique advisory services firm specialized in technology, media and telecommunications. Xona was founded in 2012 by a team of seasoned technologists and startup founders, managing directors in global ventures, and investment advisors. Drawing on its founders' cross functional expertise, Xona offers a unique multi-disciplinary integrative technology and investment advisory service to private equity and venture funds, technology corporations, as well as regulators and public sector organizations. We help our clients in pre-investment due diligence, post investment life-cycle management, and strategic technology management to develop new sources of revenue. The firm operates out of four regional hubs which include San Francisco, Paris, Dubai, and Singapore.

Xona Partners

www.xonapartners.com

advisors@xonapartners.com