Hosting Options for Utility LTE Networks

UBBA
Utility Broadband Alliance

The purpose of this paper is to provide an introduction to the major three core architectures for Utility Private LTE Networks: On-Premise Utility Hosted, On-Premise Vendor Hosted, and Off Premise Vendor Hosted. There are slight variations, but these are the three major types. The paper provides a list of the critical attributes for the reader to consider and compares.

Thank You to the Members of the UBBA
Cybersecurity & Technology Architecture Working Group

Utility Network Architecture Options

Utilities, like several other industry verticals, need to decide whether to host the network On-Premise or remote. They also need to decide whether to provide the network functionality themselves or contract the build/operate out to external vendors that provide cloud/ hosting services. There is no one solution that fits every situation. The utility must determine their requirements and priorities. Once those are crafted, only then can they make an informed decision on the architecture that works best for their environment. The attributes presented here are designed to provide “food-for-thought” and guidance towards the decision process.

Table of Contents

Network Architecture Overview

To understand the high-level network deployment options, let us first get a basic understanding of the multiple layers involved within a utility network (see Fig. 1). At the lowest layer is the utility asset network, consisting of generation, transmission, distribution, and the variety of endpoints that may attach to the grid. These could be smart meters at residential neighborhoods and businesses, EV charging stations at malls, and other assets like bucket trucks, transformers, substations etc., which may require connectivity and coverage. All of these are "assets" that may belong to the utility.

The next two layers above this layer is the cellular connectivity layer. This can be broken itself into three parts – the Radio Access Network (RAN) consisting of macro towers, small cells, remote radio units (RRUs) and other communications gear that is responsible for radiating coverage in a certain frequency to a geographical area or a specific site. All this RAN gear connects back to the Core Network Layer via a back-haul fiber network (noted in the picture).

Having an understanding of the various layers, a utility can now easily compare and contrast the various deployment options for the architecture. The rest of the white paper compares and contrasts these options along certain key axis (like scalability, latency, risk, geographic flexibility, product availability), allowing the reader to make an informed decision on which options might be most appropriate for them.

The Core Network Layer consists of the 4G/LTE or 5G Core, responsible for providing voice, video, push to talk and other value-added services to endpoints in the asset layer. It also contains breakout gateways to connect the endpoints to the public internet (if so configured). Within the same layer, we have the OA&M (Operations, Administration and Management) front-end for allowing the operator to monitor and control the performance of the network.

The final layer is the Application layer, which assumes some form of connectivity to the assets in the asset layer (by virtue of the Network layers below), and can provide data-collection, analytics, grid management, and vertical applications such as mobile workforce management (MWM), and vegetation management, etc. These applications can be developed by the utility itself or could be third party applications that run in the cloud.

Typical Utility Network Architecture

Fig. 1

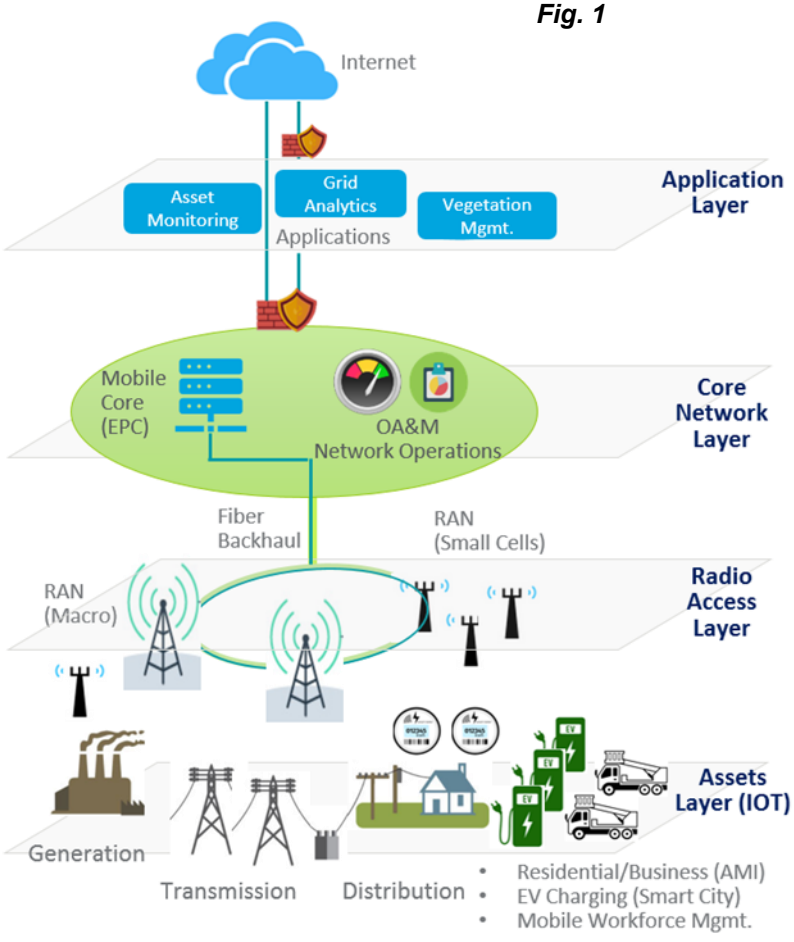

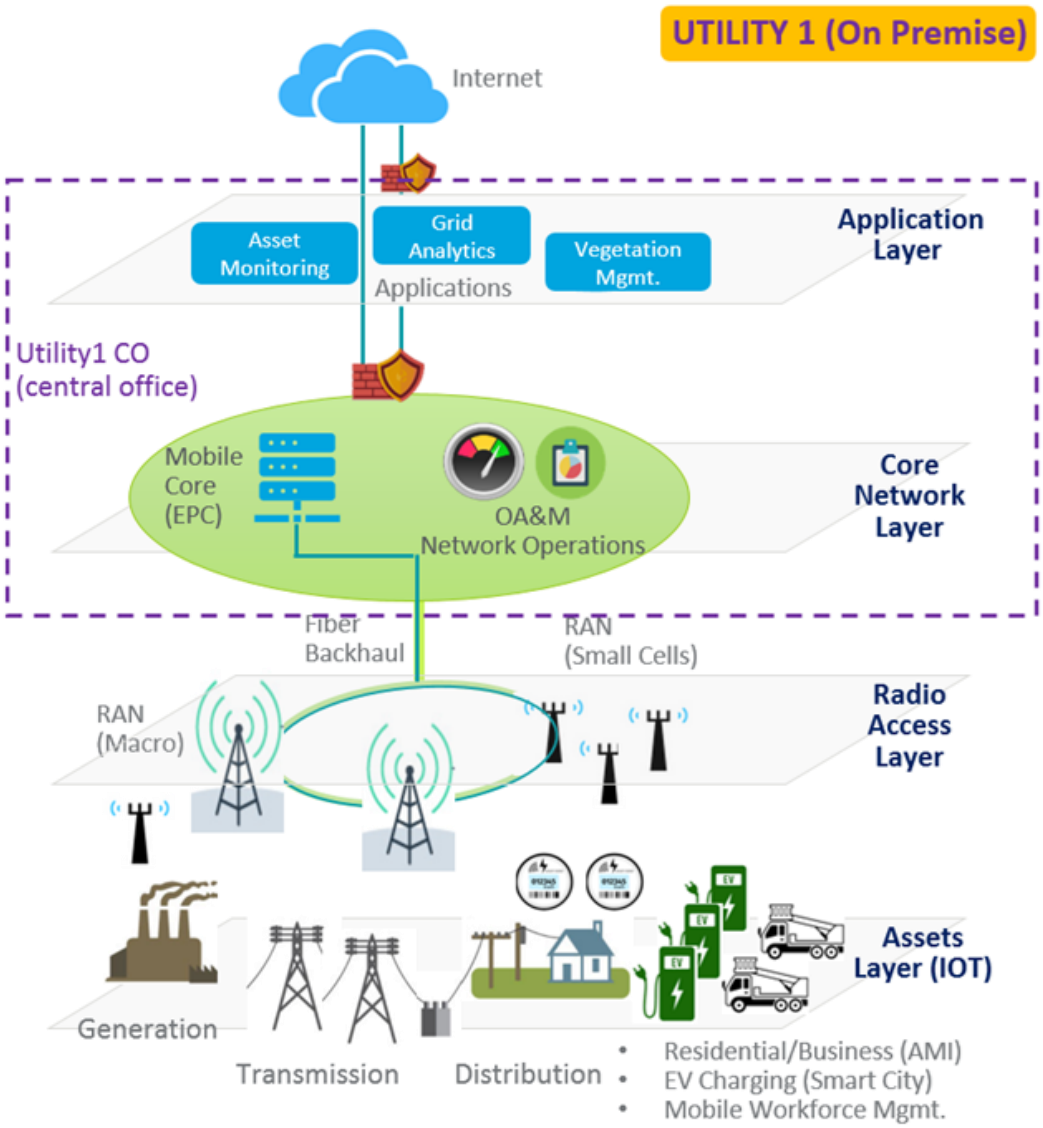

Fig. 2

On-Premise Utility Hosted

In this option, the utility acquires the hardware and software for running the mobile core network and the area shown in Fig. 2 in the dotted rectangle, runs on the utility's premises. The on-premises option is useful for utilities that have already invested in hardware and are looking to have total control over the solution. The utility would need to develop expertise in operating a telecom network, which in all likelihood is outside of its core competency today.

The role of this networking team morphs as the project progresses from Plan to Build to Maintain. While the utilities ramp-up their internal staff and skillsets, it is recommended that they lean on the Original Equipment Manufacturer (OEM) vendor and/or telecom implementation experts (systems integrators) to get through the Plan and Build phases and take over control as the network moves into the Maintain stage. We recommend that the utility staff that will maintain the network long term be a part of the plan and build phases of the project to learn about the systems being installed.

Network Hosting Options - Cloud

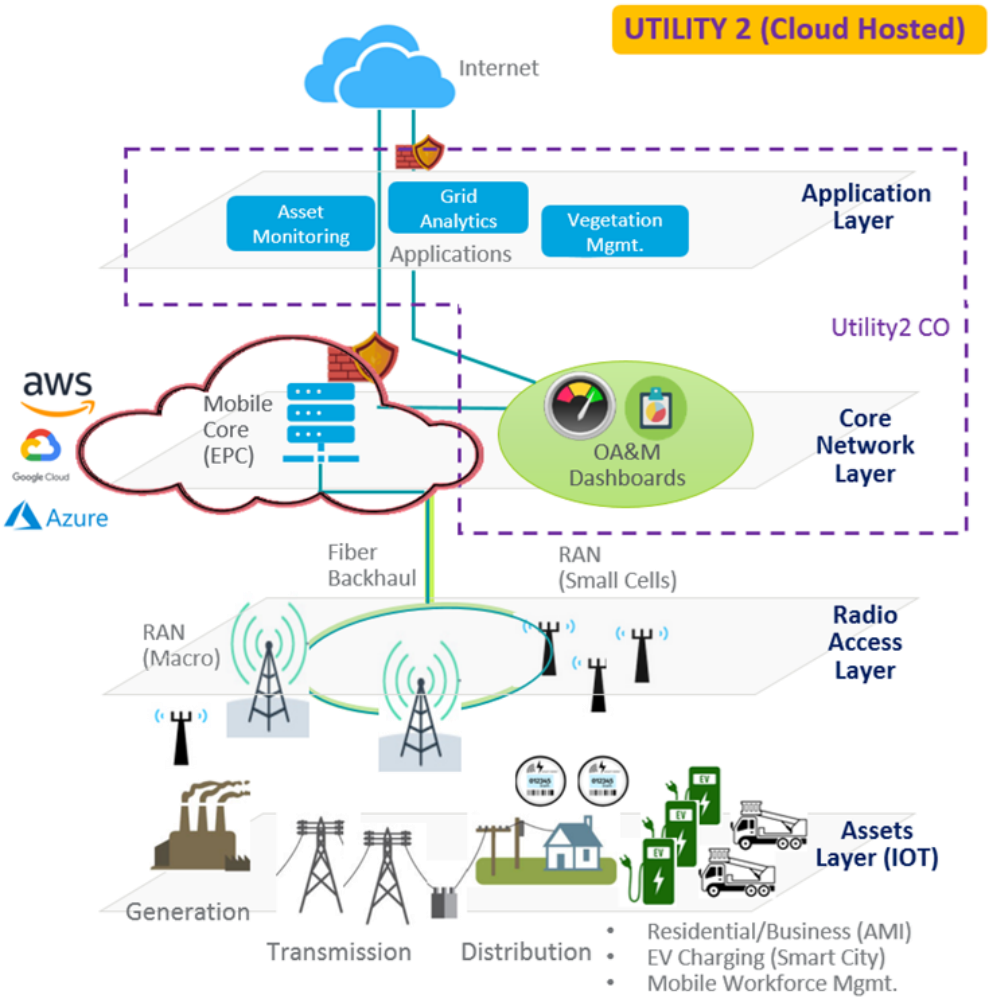

Fig. 3

On-Premise Vendor Hosted

The picture for vendor hosted option looks no different from the Fig.2, except for the business relationship. Rather than utility owning and operating the hardware and software required for operating the mobile core network, a vendor provides a functioning core network (including the hardware) to reside at the utility's premises. The vendor takes responsibility for the hardware, upgrades, and patches, and ensures the mobile core maintains an uptime that meets or exceeds the agreed to Service Level Agreement (SLA). The applications reside at the utility premises. The utility manages the service provisioning (moves, adds, drops) whereas the Vendor manages the scaling and availability of the platform.

Vendor hosted option reduces the responsibility of the utility in terms of networking setup and figuring out the CPU/ memory / disk options for the required hardware. The utilities can focus their staff on building the right applications for managing their energy network, their field services, metering, etc.

Remote Hosted - Cloud Providers

Remote hosted option hosts the mobile core network in data centers of the cloud providers (like Azure, Amazon Web Services AWS, and Google Cloud) outside of the utility premises. Both AWS and Microsoft, for instance, now have specific telco offerings (e.g., Azure for operators and AWS for private wireless 4G/5G), where a functioning mobile core network is available to be spun-up on demand. The dashboards for near real-time control and decision making would continue to reside in utility premises, along with any applications purpose built for the specific utility. See Fig. 3.

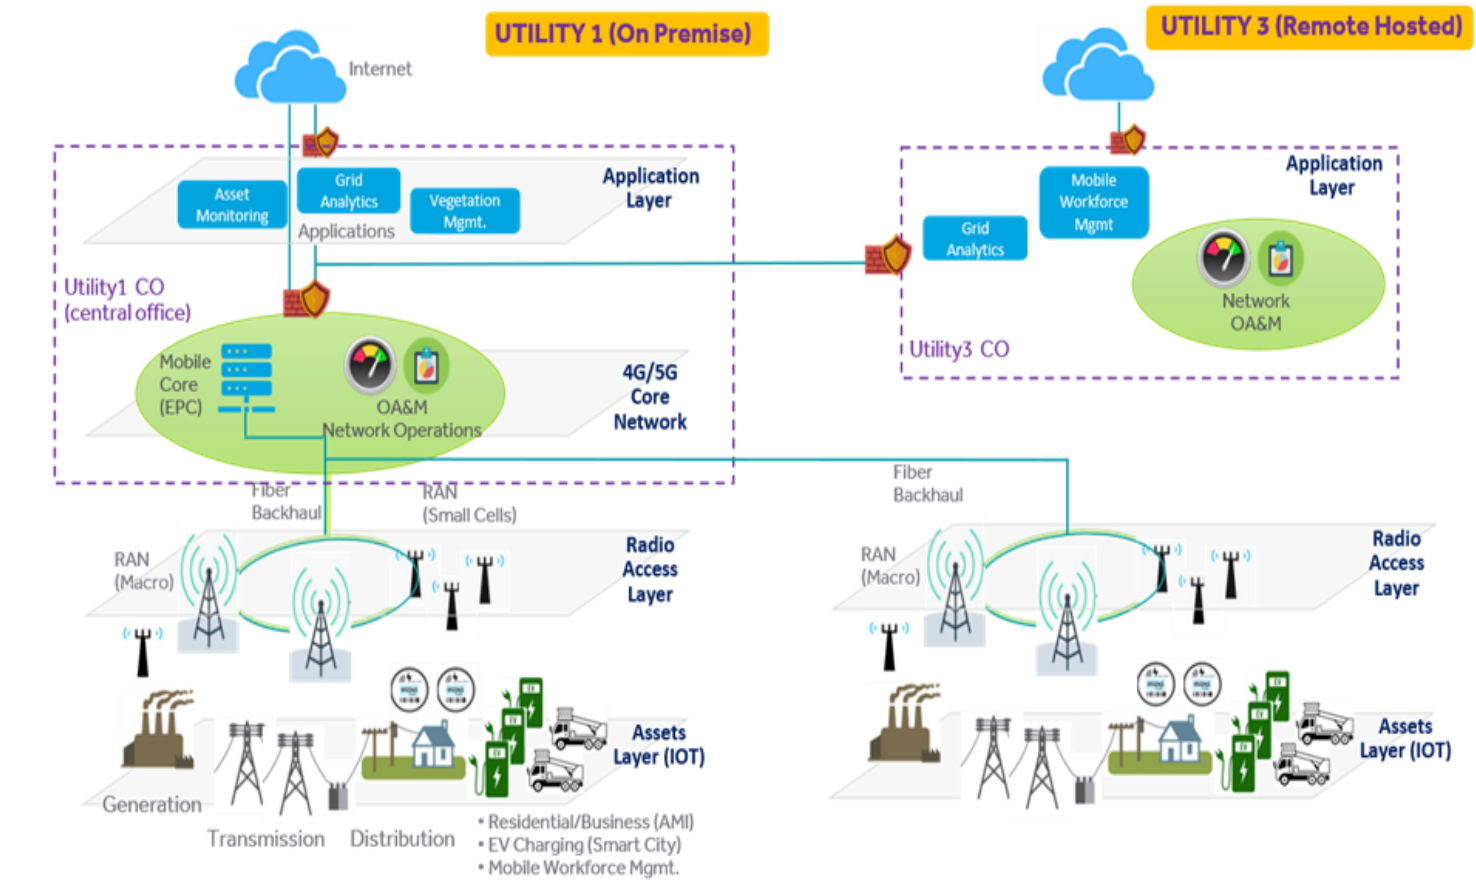

Fig. 4

Remote Hosted - Larger Utilities

Remote hosted option hosts data in data centers of larger utilities. For utilities, given their geographical boundaries would be a feasible option as large utility acting as a remote hosting or backhaul site for smaller utilities. This is a win-win for both utilities in terms of potential cost savings and of reliability and security expectations compared to dealing with a third party vendor.

Another advantage of having a larger utility host the core network is that they may be able to assist the new utility with potential challenges with LTE implementation. Most SIM card vendors for example, are used to working with Tier 1 and 2 LTE carriers and selling SIM cards in the 100's of thousands. Procuring your own SIM Card and SIM profile for less than 50,000 SIMs may be extremely challenging.

The core hosting utility has already worked through these challenges. setting up redundancy in their networks. In this model the RAN part might be built/expanded by the smaller utility, but it is backhauled to a shared core running at a larger utility.

In Fig. 4, Utility one is sharing its mobile core network and extending it to the RAN infrastructure of utility three. For utility three, this is no different than a cloud-hosted option, except that the remote cloud is being provided by utility one. A utility considering implementing LTE may do it for many reasons, but utility grade reliability and security is likely a big one. Working with a larger utility that shares a similar mindset may help offset frustration and misunderstanding.

ATTRIBUTE REVIEW BY ARCHITECTURE

On-Premise Utility Hosted

Scalability

There are two basic types of scaling in cloud computing: vertical and horizontal scaling.

To scale vertically (scaling up or down) a utility could add or subtract capacity to an existing cloud server by upgrading memory (RAM), storage or processing power (CPU). Usually this means that the scaling has an upper limit based on the capacity of the server or machine being scaled. This often requires downtime to implement but it can reduce data center footprint.

To scale horizontally (scaling in or out), a utility could add more resources like servers to the system and spread the workload across them, which in turn increases performance and storage capacity. This increases availability as it requires minimal down-time, but adds additional servers increases the footprint in the data center.

For the on-prem utility hosted model, the utility is responsible for maintaining infrastructure and determining how to scale when needed. It would need processes and procedures to monitor the capacity of the cloud infrastructure and identify thresholds for scaling (either vertically or horizontally) to support network growth. A third party configuration management service or tool could also help manage and optimize scaling.

Latency

In this scenario the latency is limited to the latency inherent in the LTE network and is low enough to meet all use cases.

Customization

The infrastructure architected and managed on-premises is limited only by the budget and the amount of floorspace available in which to build. Things to take into consideration that may be out of a utility's control include the power being sourced to the building (feeds from multiple utility substations are ideal for redundancy reasons) and location of the building itself (what may have been ideal placement for a building's initial purpose may not be so for use as a data center when it comes to security, flood plain, or hardware delivery).

Security

Security is typically considered in two categories, physical security, and logical security. Perhaps the biggest risk to any security plan is the human element so developing specific policies and training around them is necessary to successfully protect the network.

⇒ Physical Security - In the on-prem utility hosted model, the LTE core(s) can be housed in existing utility datacenter(s) and should therefore be subject to existing policies around access to and surveillance of that facility. A utility may want to consider whether additional physical barriers around the LTE core hardware to further limit access to only those trained and qualified to manage / maintain the equipment.

⇒ Logical Security / Remote Access - In the on-prem utility hosted model, the connections outside of the utility network necessary to manage the LTE network are minimized and all data traffic can stay inside the secure utility network. An IPsec tunnel and specific customer managed authentication will need to be implemented for OEM remote access to the LTE core for support. Each OEM will have an established process for secure connectivity and authentication of individuals accessing the core remotely. If a utility has commercial carrier roaming as an implemented use case, then secure VPN connectivity to an IPX vendor at two of its data centers will also be necessary.

⇒ Connectivity – Direct or Internet -In this case secure connectivity via VPN / IPsec over IP is a viable and economical option versus direct connectivity.

Continued

On-Premise Utility Hosted - continued

Technology Upgrade

The technology upgrade ability in any of these scenarios is made significantly easier with the use cloud-based infrastructure rather than legacy OEM native hardware. The former are simply software upgrades with minimal, if any, downtime while the latter requires a forklift upgrade of both hardware and software with significant downtime.

Financial

There are several factors to consider when looking at financials when deciding on an architecture. First, a utility must decide if they have Capital or O&M dollars to spend. A utility owned core could maximize capital spend whereas a hosted solution could maximize O&M spend. The Accounting Department will need to provide guidance on what is considered capital vs. O&M.

Second, a utility must consider if it wants to spend dollars upfront or if it wants to delay costs and spread-out expenditures. A utility owned solution would require a sizable outlay initially whereas a hosted solution cost could be spread over years.

Third, a utility needs to consider other factors around the timing of the spend. Does the utility have a rate case, or a specific budget spend that requires the timing to be completed by a specific deadline or delayed for a specific amount of time? A utility owned, and hosted solution requires a substantial outlay of cash initially with the ability to capitalize that spend.

Risk

Risk has many forms. The most common are reliability, resources, and financial. Reliability can be a risk if the utility does not have the technical expertise to build and maintain a core efficiently. Even with initial help from the vendor or an outside engineering firm, a utility could struggle with the maintenance. If the core is not configured properly or well-maintained, reliability could suffer.

Resources are another concern. If a utility doesn't have the technical staffing, or the space in its data center, the utility could compromise a quality design and short-cut their requirements.

Financially, a utility needs to ensure that it has the funding to support the size and scope of its design. Underestimating or overestimating the number of users on the system could impact the spending model and leave a utility with not enough funding to adequately cover the project.

A utility owned solution will give greater control, but will require the utility to assume the risk for the core.

Management

Management of the core is dependent on the level of control needed over the system. For example, controlling software upgrades may be crucial to coordinating outages with critical users. A vendor hosted solution may not be as accommodating when coordinating outages.

Alarm management is another area of control that may be desirable. If the utility has the resources in their Network Operations Center (NOC), they may want to control and closely track alarms because they want to provide their own first level of troubleshooting. A utility may be able to provide alarm resolution faster than a hosted solution if they have the resources.

Activation of new services and adding users may be another reason for managing its own core vs. a hosted solution. Again, a quicker response would be a benefit of a utility hosted solution. A utility hosted solution provides greater control, but requires more resources.

Redundancy

Redundancy of the LTE core is a must when implementing highly available network. Generally, there are two standard implementations of redundancy, locally and geographically. Local redundancy or high availability is achieved in the configuration and architecture of the core itself. The core hardware and software are configured in such a way that multiple instances of CNF/VNF (Container/Virtual Network Functions)are running simultaneously and none of them are on the same physical host. This allows high availability in the event of a hardware or software failure. Geographic redundancy

Continued

On-Premise Utility Hosted - continued

further increases availability of the LTE network because a utility has two LTE cores supporting the network in two diverse locations.

In the event of a larger failure or data center failure, the redundant location can assume control of the cell sites and/or core functions depending on the nature of the failure. In the on-prem model, utilities can take advantage of their existing data center infrastructure redundancy and connectivity to house the core nodes.

Product Availability

Most OEMs' LTE core solutions are designed to support quantities of devices, cell sites and data throughput and many have several options depending on those quantities. If the utility network is small geographically, in number of cell sites and number of devices then there is likely a core solution with a small footprint, limited capacity, lower in cost, and relatively easy to deploy. Likewise, if the utility network is large in geography, cell sites, devices, and anticipated throughput then there is a much larger solution with a larger footprint, more complexity and more expensive. Utilities should understand their needs and work with the OEMs to identify a solution that fits their current and future plans. In the on-prem model any LTE core solution is viable.

Speed to Turn-up

Speed to market is slowest in this scenario because the hardware does not exist and must be procured, installed, and commissioned before being put into service.

Space for Equipment

Accommodation for production hardware and the storage of additional equipment and parts is limited by the space available to a customer. The installation and potential for future growth must be accounted for and included in facility planning to account for power and cooling capacity.

Outages

Planned outages are under the direct control of the utility which can play a major role in the reliability of the system and the scheduling of the maintenance. It also requires resources in this configuration because the utility carries the responsibility for addressing unplanned outages.

ATTRIBUTE REVIEW BY ARCHITECTURE

On-Premise Vendor Hosted

Scalability

In this option, the utility uses a third-party cloud provider such as Azure, or AWS, or Google Cloud to deploy on-prem cloud infrastructure to support the deployment of the LTE EPC. Scale could again be achieved vertically or horizontally ,but the third-party cloud providers also have vast hardware and software resources in place off-prem that could allow for even more rapid scaling.

Latency

If scaling is confined on prem, the latency is limited to the latency inherent in the LTE network and is low enough to meet all current use cases. However, if scaling to meet demand is off-prem, then additional latency is introduced. Because third-party data centers can be located anywhere in the world utilities should take time to identify and possibly even require data center capacity that minimizes latency.

Customization

A vendor hosting a solution on-prem means providing power and space to their equipment, so any customization is limited to the hardware platform agreed upon and the application(s) running on it. The level of hosting could be anything from a complete solution wherein the vendor provides all necessary hardware which would run within a customer's space to more of a shared-hosting environment in which the vendor might provide the servers on which an application runs but are connected via the customer's existing network.

Security

Security is typically considered in two categories, physical security, and logical security. Perhaps the biggest risk to any security plan is the human element so developing specific policies and training around them is necessary to successfully protect the network.

⇒ Physical Security - In the on-premise vendor hosted model, like the utility hosted model, the LTE core(s) can be housed in existing utility datacenter(s) and should therefore be subject to existing policies around access to and surveillance of that facility. A utility may want to consider whether additional physical barriers around the LTE core hardware to further limit access to only those trained and qualified to manage / maintain the equipment.

⇒ Logical Security / Remote Access - In the on-prem vendor hosted model, the connections outside of the utility network necessary to manage the LTE network are like the utility hosted model in that the OEM will need remote access for support and secure connectivity to an IPX vendor is needed for commercial carrier roaming. In this model however the vendor is also managing day-to-day operations of the network. Typically, this OEM managed service organization is separate from the OEM customer support organization and will likely require a separate secure connection to a different part of the OEM network. The OEM will have policies and procedures around connectivity and authentication for these managed services as well.

⇒ Connectivity – Direct or Internet - In this case secure connectivity via VPN / IPsec over IP is a viable and economical option versus direct connectivity.

Technology Upgrade

The technology upgrade ability in any of these scenarios is made significantly easier with the use cloud-based infrastructure rather than legacy OEM native hardware. The former are simply software upgrades with minimal, if any, downtime while the latter requires a forklift upgrade of both hardware and software with significant downtime.

Continued

On-Premise Vendor Hosted (continued)

Financial

On-premise vendor hosted solution is a hybrid mix between ownership and cloud environment. Traditionally the software licensed is procured directly and the accessibility is found in the cloud via a server that is hosted and managed by the vendor. Cost could be more, as the cloud environment is private and could proliferate additional cost with support and integration with a virtual private cloud hosting provider. There is a larger upfront expense with the software being paid upfront ultimately absorbing the cost of IT infrastructure, maintenance, and administrative staff. With procurement of the software license, a utility has ownership of the software and can be used indefinitely. The solution could potentially run on offline systems. It is a great way for utilities to remove the inconvenience and expense associated with managing servers, security, upgrades, and the like. The cost of ownership doesn't have the potential volatile variability that an off-premise solution has, where the total cost of ownership may go up after a couple of years.

Risk

The burden of risk for managing the core shifts to the vendor. The utility does not have the risk of system failure or the risk of finding the correct resources to properly maintain the core. Latency risk is also minimized with the equipment residing on site. The risk is all associated with the vendor which could be important if the utility has a low risk tolerance.

Management

While the utility maintains ownership of the software database, a vendor would be responsible for software maintenance and management, including things like backup and upgrades, thus freeing up resources and potential required skill set levels. On-premise vendor hosted usually includes benefits of fast performance, scalability of the cloud storage, data security and scheduled backups. It takes more of a managed approach on the vendor side, taking a lot of the risk for data security, integrity, and storage out of the scope of work for the utility. Before a utility goes throwing the on-premise servers away, some things to think about: maintenance fees may rise over time, customization to the software could result in additional charges, and vendors access to customer data. A utility should discuss these topics with the vendors being considered. The burden of monitoring, maintaining, and troubleshooting now lie with the vendor.

Redundancy

Redundancy the LTE core is a must when implementing highly available network. Generally, there are two standard implementations of redundancy, locally and geographically. Local redundancy or high availability is achieved in the configuration and architecture of the core itself.

The core hardware and software are configured in such a way that multiple instances of CNF/VNF are running simultaneously and none of them are on the same physical host. This allows high availability in the event of a hardware or software failure.

Geographic redundancy further increases availability of the LTE network because it has two LTE cores supporting the network in two diverse locations. In the event of a larger failure or data center failure, the redundant location can assume control of the cell sites and/or core functions depending on the nature of the failure. In the on-prem model, utilities can take advantage of their existing data center infrastructure redundancy and connectivity to house the core nodes.

Product Availability

Most OEMs' LTE core solutions are designed to support quantities of devices, cell sites and data throughput and many have several options depending on those quantities. If the utility network is small geographically, in number of cell sites and number of devices then there is likely a core solution with a small footprint, limited capacity, lower in cost and is relatively easy to deploy. Likewise, if the utility network is large in geography, cell sites, devices, and anticipated throughput then there is a much larger solution with a larger footprint, more complexity and more

Continued

On-Premise Vendor Hosted - continued

expensive. Utilities should understand their needs and work with the OEMs to identify a solution that fits their current and future plans. In the on-prem model any LTE core solution is viable.

Speed to Turn-up

Speed to market in this solution is comparable to the on-prem solution for the same reasons. If the third party vendor has hardware on hand that would improve the time to market.

Space for Equipment

Accommodation for production hardware and the storage of additional equipment and parts is limited by the space available to a customer. The installation and potential for future growth must be accounted for and included in facility planning to account for power and cooling capacity.

Outages

This architecture would be the same as the Off Premise scenario. See Off Premise (Vendor Hosted) section.

ATTRIBUTE REVIEW BY ARCHITECTURE

Off-Premise Vendor or Utility Hosted

Scalability

In the off-premise vendor solution, the utility uses a third-party cloud provider such as Azure, or AWS or Google Cloud to host the deployment of the LTE EPC off-prem. The utility can leverage the existing massive scale of these providers to deploy the LTE EPC and scale easily and rapidly to accommodate capacity. The vendor is again responsible for maintaining the infrastructure, monitoring capacity, and adding as necessary reducing operational load on the utility.

In the off-premises utility-hosted model, the utility uses another utility that already has an LTE core to host the deployment of the LTE EPC. The scalability of this solution will depend on the host utility's willingness and ability to meet your future UE and RAN site demands.

Latency

In this solution, similar to the on-prem vendor hosted, additional latency is introduced into the network. A way to mitigate this is with a hybrid vendor hosted solution both on and off-prem. Leveraging the CUPS (Control and User-Plane Separation) architecture introduced in 3GPP Rel. 14, the user plane functions can remain on-prem while the control plane functions can go off-prem.

Customization

Software-as-a-Service (SaaS) – using software, typically via web interface, where all accompanying services are hosted and maintained by another entity. Examples include WebEx & Google Drive. Benefits are quick-to-launch services and short-term project use; limitations include vendor lock-in where the data may not be easy to export, potential integration issues if non-standard protocols are used, and managing data security and compliance. Little to no customization offered and application usage is limited to vendor offerings for functionality and performance.

Platform as a Service (PaaS) – provides a platform for software development without having to consider servers, storage, & operating system and their related updates. Benefits include simpler development of highly available apps and accessibility to numerous developers using the same tools. Limitations are similar to SaaS in that integration with existing systems can be complex and vendor lock-in are in play. Customization typically limited to the application layer and integration with on-prem solution may require complex configuration.

Infrastructure-as-a-Service (IaaS) – this is what most people know as "the cloud", wherein servers, storage, and networking are hosted such as in Microsoft Azure or Amazon's AWS, and needed resources are purchased and available on-demand. Virtual servers, databases, and their associated storage can be spun-up and shut down at will without consideration of the underlying hardware. Clients can architect and control a computing model similar to what they run on-prem with the caveat being they are on shared hardware, so threats cannot be secured to the level of completeness possible with an on-premises solution (communication routes, host system vulnerabilities, people with access to the hardware, etc.).

Security

Security is typically considered in two categories, physical security, and logical security. Perhaps the biggest risk to any security plan is the human element so developing specific policies and training around them is necessary to successfully protect the network.

⇒ Physical Security - In the off-prem vendor or utility hosted model, the LTE Core is hosted in the vendor or the host utility data center. The utility will need to ensure that physical security policies and procedures implemented by the vendor meet or exceed the utility standards. Likely the two should be very similar but some due diligence in understanding how physical security is handled is a must.

⇒ Logical Security / Remote Access - In the off-prem vendor or utility hosted model, there are two options for deploying the LTE core architecture. In the first option, the entire LTE core is off-prem and therefore no management connectivity is necessary.

Continued

Off-Premise Vendor or Utility Hosted (continued)

⇒ However a secure connection to the OEM is still needed because all traffic will need to go off-prem to the LTE core and then return to the utility network to be routed to its final destination. In this scenario a direct connection to the OEM facilities may be a more secure option than the internet. Having data leave the utility network may not be an ideal option but the LTE core is flexible and the data carrying node (PDN Gateway) could be deployed on the utility premises which eliminates security concerns associated with data leaving the utility network. Connectivity to the rest of the off-prem core is still necessary but is limited to signaling and OEM traffic.

⇒ Connectivity – Direct or Internet - If the utility chooses to host its entire LTE core off-prem then it is advisable to have direct connectivity to the vendor where possible because all data traffic is leaving the utility network. If the utility chooses to host the PDN Gateway on-prem then internet becomes an option. From a robustness standpoint, most utilities will choose to connect directly to the LTE core via a dedicated connection.

Technology Upgrade

The technology upgrade ability in any of these scenarios is made significantly easier with the use cloud-based infrastructure rather than legacy OEM native hardware. The former are simply software upgrades with minimal, if any, downtime while the latter requires a forklift upgrade of both hardware and software with significant downtime.

Financial

Rather than spending capital on implementing an EPC, with a cloud hosted model, either a vendor or a host utility has purchased and runs the EPC architecture. To help cover some of these costs, there will be some type of fee to cover the subscription and maintenance/support costs. These fees are likely to be outlined to cover a predetermined period of time with payments made incrementally, (typically monthly). The end user does not have to have any qualified IT staff onsite to participate and all of the updates to the infrastructure and applications are included in fees. As part of the contract, there may be some one-time fees to stand up the server, but generally speaking the upfront capital costs for vendor or utility hosted model are much lower than an on-premise model. To that end, the fees for cloud hosting are not considered a capital expenditure, which means the costs are primarily considered an operating expense.

This could potentially help with budgeting and approval for projects. The treatment of these expenses between O&M and Capital would be determined by the utility's Accounting Department. The treatment of these expenditures and available funding will determine how well it fits into the utility's plans and budget.

Risk

For the most part, the risk here would be very similar if not less than the risk of an On-Premise Vendor Hosted solution.

Management

All the access, performance, updates, security, and encryption, as well as system backups are part of the vendor or utility hosted fees. It eliminates the need for a lot of in-house expertise and could save on quite a bit of time and space. Without the need to have a physical location to store the server(s), a utility can look to allocate that space for other potential projects or resources. In addition, the need to upgrade or change server size would be completely on the vendor or host utility thus taking away the need for maintenance and management of storage and computing power. There is assumed no performance degradation as the data center is always being monitored and upgraded to deliver maximum performance. Access to the cloud is easily managed with customized user fields and can be accessed (with appropriate credentials). It is important to have direct connections for both robustness and security purposes. In the current environment it would be remiss to not discuss the importance of security and securing data, thus cloud hosted models traditionally include options for the frequency of data backups, offsite data storage, and data testing not typically considered when looking at on-premise solutions.

This scenario is very similar to the On-Premise model, but there is no control over hardware or

Continued

Off-Premise Vendor or Utility Hosted - continued

security and access is remote. Security, monitoring, and upgrades all in the hands of the vendor or host utility. Because remote links are required, there is more risk of losing connectivity so redundant connections are required. An advantage is that the utility does not need to provide the resources for managing the core or the space.

Redundancy

Redundancy of the LTE core is a must when implementing highly available networks. Generally, there are two standard implementations of redundancy, locally and geographically. Local redundancy or high availability is achieved in the configuration and architecture of the core itself. The core hardware and software are configured in such a way that multiple instances of CNF/VNF are running simultaneously and none of them are on the same physical host. This allows high availability in the event of a hardware or software failure. Geographic redundancy further increases availability of the LTE network because a utility has two LTE cores supporting the network in two diverse locations. In the event of a larger failure or data center failure, the redundant location can assume control of the cell sites and/or core functions depending on the nature of the failure. In the off-prem model, both types of redundancy can still be achieved, however, geographically the utility will be using vendor locations. This will impact connectivity, security and core architecture and this will need to be investigated fully.

Product Availability

Most OEMs' LTE core solutions are designed to support quantities of devices, cell sites and data throughput and many have several options depending on those quantities. If the utility network is small geographically, in number of cell sites and number of devices then there is likely a core solution with a small footprint, limited capacity, lower in cost, and is relatively easy to deploy. Likewise, if the utility network is large in geography, cell sites, devices, and anticipated throughput then there is a much larger solution with a larger footprint, more complexity and more expensive. Utilities should understand their needs and work with the OEMs to identify a solution that fits their current and future plans. In the off-prem model some LTE core solutions that are available for deployment may not be available for OEM hosting and some OEMs may not offer hosted solutions as an option. Because of this is a multi-vendor solution may come into the picture which would further complicate the deployment and operation of the LTE network.

Speed to Turn-up

Speed to market in this solution is the fastest as much of the cloud hardware infrastructure is already in place. If there is a desire to go to market as fast as possible, some trade off in latency can be made to deploy the full EPC off-prem while deploying the on-prem PGW and then migrating to that solution.

Space for Equipment

Space needed for hardware is not a concern for cloud-hosted solutions

Outages

When the utility is hosting and controlling its entire LTE network, it will be responsible for establishing its own planned and unplanned outage communication and response processes. When having a third party host the core, it is essential that the utility take the time to document and agree upon outage expectations. Having a clear expectation on the communication of planned outages between the third party that is hosting the LTE core and the utility is critical to avoid frustration and dis-satisfaction. Depending on the end user applications using the LTE network, the acceptable time for outages and upgrades may be different than what the third party is expecting. For example, SCADA users may want outages during the daytime so that they can potentially have personal ready to respond. This may differ from the third party which may be used to scheduling planned outages or upgrades during the middle of the night. In addition to establishing clear expectations for the timing of outages, it is important to agree upon how much advanced notice is required for planned upgrades or outages. Some of the LTE network end-users may require weeks of advanced notice due to regulatory concerns especially with nuclear

Off-Premise Vendor or Utility Hosted - continued

facilities. If a third party hosts the core, make sure that advanced notice timeframes for outages are clearly documented in the agreement.

In addition to planned outages, the communication of unexpected outages is critical. Setting clear requirements on dashboards, alarm management, and verbal and electronic communication is essential. The utility's current telecom control center has likely already established a set process for communication with management or end users on outages. It is important to think through and document communication expectations for both communication type, timeliness, and cadence for the third party that is hosting the core.

In addition to human communication, many utilities may also want visibility into real time alarming that is occurring on the LTE hosted core. Having alarms forwarded to the network management platform will help the LTE control center have real time perspective and notice of LTE core issues. Dashboards can also be useful showing critical LTE Core alarm status. If a utility is working with a third-party to host their LTE core, make sure that the communication of planned and unplanned outages is thought through and documented.

-END-

APPENDIX

SUMMARY OF ATTRIBUTES FOR EACH OF THE ARCHITECTURES

Architecture				
Attributes	On-Premise Utility Hosted	On-Premise Vendor Hosted	Off-Premise Vendor/Utility Hosted	Comments
Scalability				
Latency				
Customization				
Security/ Remote Access				
Technology Upgrade				
Financial Risk				
Management				
Redundancy/ Geographic Flexibility				
Product Availability				
Speed to Turn-up				
Space for equip				
Outages				

Table 1: summary of those attributes and is designed to help the reader determine which architecture is best for them based on their situation in each of the attributes. The reader can us the table as a score card to help guide them to the best solution for their needs.