WHEN COVERAGE ISN'T ENOUGH:

Designing for Service in Utility Private LTE Networks

CRISTINE KOROWAJCZUK
CCO, CelPlan Technologies, Inc.

When people ask me about designing RF networks for utilities, I often start with a concept that many overlook: not all coverage is created equal. After discussing adaptive modulation and interference management in previous articles, let's dive into the distinction between coverage area and service area – a critical concept that can make or break your utility's private LTE deployment.

Coverage means enough signal is present for your radio to detect it and connect to a cell. But service requires more: the radio must connect, transmit, and meet the application's performance requirements – things such as throughput, bit error rate, and latency. Your field crew may have "signal bars" on their devices, but if it takes ten seconds to send telemetry or if it fails to initiate a session during a fault event, then the coverage isn't meeting your service needs.

I've seen multiple utility networks start design, and sometimes even deployment, with the wrong premise. They assumed they needed coverage everywhere, when in reality, their assets and use cases were concentrated in specific areas. This matters because utilities don't deploy networks just for connectivity – they deploy them for mission-critical applications, each with different performance profiles. To illustrate this, imagine you have cellular coverage in a building's basement. Your phone shows one bar and you can send text messages, but you can't stream video or make a reliable voice call. You have coverage, but not service for those specific applications. This brings us to an important planning consideration: how do we define where service is actually needed? When planning a private LTE network, it's common for utilities to define their Area of Interest (AoI) based on their entire service territory. But are you really trying to serve every square foot equally? A more accurate approach is to define AoI based on where communication is actually needed for each device type.

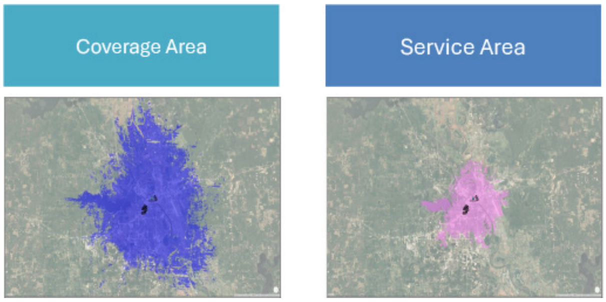

Figure 1: Comparing Coverage and Service Area

When a utility tells me they need "complete coverage," I ask them three main questions: What kind of devices and use cases are you actually serving? Where are these devices most likely located? Are they fixed or mobile? The answers to these questions often reveal an important misjudgment – the actual network AoI is significantly different from what was initially envisioned. I've worked with utilities that planned for blanket coverage across their entire service territory, only to discover that 90% of their critical assets were concentrated in 20% of the area. Typical fixed critical communication devices are usually located at substations, along feeders or distribution lines, in residential and commercial areas, or close to generation assets. Mobile devices usually include field crews and inspections (for example, drones). Even when you don't know the exact location of a device, most of the time you have a good idea of its general location or the area that they will be working on.

Figure 2: Definition of Area of Interest (AoI)

Even when you don't know the exact location of a device, most of the time you have a good idea of its general location or the area that they will be working on.

Speaking of fixed versus mobile assets, let me share something that many network designers overlook. Fixed devices have a huge design advantage: they don't move. That means you can optimize their connectivity using directional antennas and careful placement to maximize high-modulation use – and, with it, spectrum efficiency. This predictability is incredibly powerful for network optimization. When you know exactly where your substations, distribution automation equipment, and smart meters are located, you can design the network to maximize their connection at higher modulation schemes. This dramatically increases your effective cell capacity without adding more infrastructure. While omnidirectional antennas might seem simpler to deploy, directional antennas can significantly improve SNIR for specific locations. I've seen networks double their effective capacity in certain areas just by strategically deploying directional antennas for fixed assets located between cell sites. Mobile devices are a different story. They can appear anywhere in the cell, including at cell edge, which means that they often trigger adaptive modulation downgrades and causing more resources to be required to transmit the same data. If you're not accounting for this mobility behavior, especially during high-demand events like storms or blackouts, you risk oversubscribing your network when you need it the most.

And, since we introduced the cell edge topic, let's talk about a counterintuitive principle that catches many off guard. By now you are probably tired of my repeating the statement that "higher modulation schemes allow you to fit more bits per symbol," which is true but only if your Signal-to-Noise-and-Interference Ratio (SNIR) is good enough to achieve these higher schemes. As you increase cell size, SNIR drops – especially at the edges – and so does your ability to use high-order modulations. The result? More devices get downgraded to lower modulations, consuming more resources per bit and shrinking your effective cell capacity.

This is why bigger cells can mean less capacity. The critical point to understand here is that the majority of your cell's area typically operates at lower modulation schemes. As you expand your cell size to cover more geography, you're actually reducing your overall capacity. It's a trade-off: shrinking the cell or using directional antennas sometimes gives you more usable

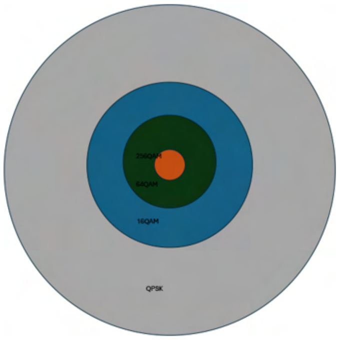

Figure 3: Cell relative area per modulation scheme (illustrative)

capacity where it you really need it. Think of it as concentric zones of service capability. Close to the tower, you have 64QAM or higher modulation, supporting high-bandwidth applications. At mid-range, you're down to 16QAM, suitable for most utility applications. And at the cell edge, you're limited to QPSK, which can only handle low-bandwidth, delay-tolerant applications.

Besides mobility, there are other things that can impact modulation schemes. Interference happens. Rain fades happen. Unexpected events happen. That's why you don't design a network for the ideal day – utilities, in particular, need to design for reality and, often times, worse case scenarios. This means leaving room for adaptive modulation to downgrade temporarily without choking the entire cell.

I always recommend planning for modulation downgrades. Temporary interference or signal variations may cause eventual modulation scheme downgrades, and it's always good to have some room to accommodate

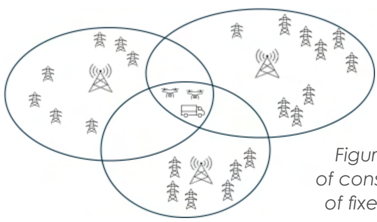

Figure 4: The importance of considering the location of fixed and mobile assets

this possible fluctuation in demand. Think of it like budgeting. If your network is running at maximum resource usage under perfect conditions, then a little interference will cause modulation to drop, and now you need more symbols to send the same data – but the resources aren't there. Something must give either latency spikes, packets drop, or users lose connectivity altogether.

During emergencies, when mobile devices flock to specific areas or interference spikes due to external factors, you need a clear prioritization strategy. Are multiple field crews going to converge on the same location? What if that location is highly prone to interference,

During emergencies, when mobile devices flock to specific areas or interference spikes due to external factors, you need a clear prioritization strategy.

for example, the area between multiple cells? What happens if not all users can be served? You need to have strategies in place to deal with "uncommon but possible" situations. Quality of Service policies that prioritize critical applications, dynamic resource allocation based on device type or user role, and backup network connections for absolutely critical communications all become essential parts of your design. Remember, in an emergency, a working low-bandwidth connection for critical SCADA commands is infinitely more valuable than a high-bandwidth connection for routine meter reads.

Designing and optimizing utilities networks requires understanding that network planning is part engineering, part strategy. It's not enough to ask, "do I have coverage?" You need to ask, "can I serve my applications, even on a bad day?" It's not enough to have an RF engineer work on your private LTE design, the whole utility team should be involved because their knowledge of the behavior of users and devices, what is essential and what we can live without, and what happens when everything fails is crucial to design and dimension a network that will properly meet the requirements.

Start with accurate AoI mapping using actual device locations, not theoretical coverage maps. Design for your critical applications first, ensuring your most important use cases have adequate service areas. Leave headroom for those inevitable modulation downgrades. Use the fixed nature of utility assets to optimize cell placement and antenna patterns for known device locations. Plan for mobile concentration points by identifying where field crews typically gather during emergencies. And always implement clear prioritization policies so you know which communications take precedence when capacity is limited.

That's the shift from designing for coverage to designing for service – and it's the foundation of a truly reliable utility LTE network. As we've explored in this series of articles, successful utility network design requires understanding the intricate relationships between modulation schemes, interference patterns, and real-world deployment constraints. By focusing on your actual Area of Interest, distinguishing between coverage and service, and considering the unique characteristics of utility applications, you can build networks that deliver reliable service where it is truly needed. ■