AT&T Business

AT&T 5G

Strategic Design Considerations for Private Cellular

TABLE OF CONTENTS

Executive Summary

The 5G cellular technology evolution is grabbing the headlines. However, other cellular developments are quietly inspiring similar enthusiasm. Private cellular is one key example.

Led by the transformation to Industry 4.0, organizations of all sizes are exploring and beginning to adopt such next-generation technologies as autonomous guided vehicles, video as a sensor and others to help drive superior operational efficiency and productivity. At the same time, current enterprise technologies are being squeezed and pushed as far as possible to the compute and network areas. In turn, these technologies and challenges are inspiring organizations to rethink how they will connect to users and things.

5G and private cellular networks are poised to play critical roles as enablers for connected sites, from a refinery to a stadium to even new Smart City constructs. Public and private sectors alike are adopting—and innovating—with private cellular. To date, however, organizations may not realize the existence of multiple private cellular solutions across different spectrum and architecture types.

This whitepaper examines essential cellular foundations in an effort to educate, demystify and offer guidance for organizations considering private cellular or other advanced wireless topologies.

Solution Designs Begin With “Where” & “What”

Taking a step back, organizations can consider a variety of wireless network options. The advent of 5G and private cellular can now be considered an alternative or complement to Wi-Fi networks. But how do you navigate a myriad of choices, solutions and considerations? Stated simply, determining the design and applicability of a wireless network begins with understanding the site and user type in context of the data creator and owner.

The first piece of the puzzle is “where”—or the site type. Sites can be interior spaces, outdoors, remote locations, dense urban environments and many others. Each requires “site-specific” solution architecture with variables influencing cost, complexity and performance, each critical in the early stages. Consider a refinery. It’s mostly outdoors, steel and pipes everywhere and requires special safety considerations. Cellular is a strong solution consideration here, but which variant? Conversely, think about an office space with abundant cubes and users working in place with fewer “endpoints” as well as lower ceilings. In this environment, Wi-Fi may likely be the right starting point—but, again, which one?

The second piece is “what”—or the user type consuming the data. At this stage, beginning with an understanding of who creates the data and who owns the data helps clarify an optimal investment and mix of networks. Imagine a person working in their cube or a baseball fan at a stadium. In both environments, cellular may be optimal in terms of performance and adding scale with Wi-Fi. With this mix, the office can maintain performance and lower relative costs, while the stadium can ensure greater user access if fan engagement is a priority. On the other hand, an employee at a vineyard with a tablet would likely be most efficient on a cellular network. The same can be said of an autonomous guided vehicle (AGV) transporting packaged items in the vineyard’s shipping facility.

Site type and user type help determine a starting point. An industrial site owner wants data to drive operational technical value. While a stadium owner might want to provide improved fan experiences and need to send data to large numbers of consumers with different devices, SIMs and abilities. There are sites with mixed user types as well. These sites can include hospitals, airports, office buildings and other sites blended with business and consumer types that may require a concentrated focus on user requirements. With an understanding of site type and user type, we can begin to weigh choices for the right wireless network solution. Throughout this paper, a starting consideration can be driven from the previous guidance and the approach illustrated in the graphic below. However, later in this whitepaper, we will introduce specific considerations to help focus a non-linear decision process. At least in the near term, it’s also worth noting it won’t be an all-or-nothing proposition as in the age-old Wi-Fi-versus-cellular debate.

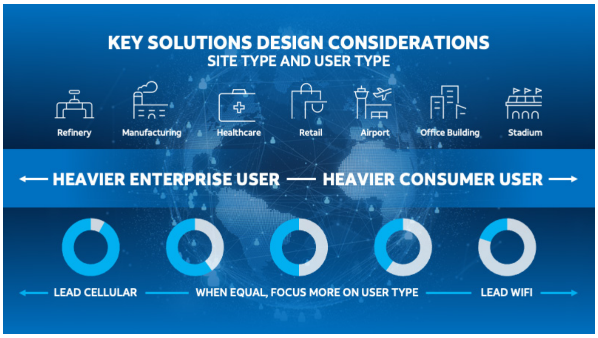

Meet The Six-Way Match

You know "where" (site type) and "what" (user type)—now comes "how." Enterprises must determine how devices will connect to the network, how much spectrum is available, how to assure transport, location, backhaul, access and how data gets stored.

Beginning with an appropriate network can address operational, business, security or efficiency problems, preventing issues down the road. Think about how data flows between devices and content as shown in the "Six-Way Match."

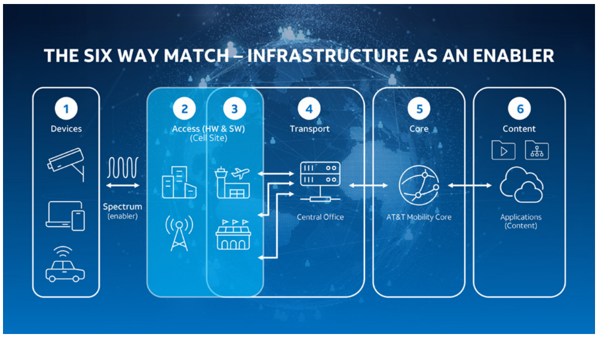

The Six-Way Match reflects the data flow or cellular infrastructure to help you build the correct solutions for your business.

Examining each of these elements will help us provide an optimized network solution that balances cost, complexity and performance.

Device Requirements & Availability

Device requirements are constantly evolving. In every case and for every solution, device type, capabilities and availability are key considerations. Private cellular networks offer an important value proposition. Devices will only work on the network and only those devices provisioned by the team can gain network access. This is an important benefit for network security when devices operate within the customer premises yet need to work in "the wild" as well. In this case, leveraging the existing public cellular network with existing devices or dual SIMs may be a practical way to achieve ubiquitous connectivity nationwide. In hybrid models, devices may be able to connect to both but that may limit device management and security in the wild.

Billions of devices are compatible with Wi-Fi and 4G LTE. When compatible 5G devices are manufactured at scale, costs will decrease. Yet, as customers move to 5G and private cellular networks, device cost and availability will become additional considerations.

While currently limited, the 5G and CBRS (Citizens Broadband Radio Service), a new spectrum in the US enabling new private cellular constructs) device portfolios will proliferate as technology and manufacturing cycles accelerate with increased interest and deployments. In general, cellular devices for these types of networks are based on the same basic technology as public cellular devices. However, bringing on a diverse selection of cellular devices in new or little-used spectrum frequencies requires time to create and manufacture the chipsets at scale. With the associated limited device availability, it's important to plan usage prioritization and implementation in conjunction with site builds.

The State Of Spectrum

Spectrum is the essential lifeblood of wireless connectivity—a vital component that makes a network wireless. Without it, nothing happens. Wireless communication occurs over a range of frequencies, and those frequency ranges are not created equal. We'll compare different ranges then take a deeper look at differences between unlicensed, licensed and shared spectrum.

The first range to consider are low-frequency bands (typically in the sub 1GHz range). Given their physical capabilities, these signals travel farthest, allowing for the greatest coverage that also can penetrate obstructions. Low-band capacity makes it ideally suited for low data-rate sensors or simple comms like push-to-talk or text messaging. When low-band frequencies are used with higher capacity payloads, users should expect longer transmit times given the smaller "pipe" size relative to others.

At the other end of the scale are high-frequency ranges (including mmWave in the 24GHz range and higher) that travel shorter distances with reduced abilities to penetrate obstructions as the frequencies increase. These high frequencies allow for the largest amount of capacity. That means, it's easier to move large workloads that might feature video or a virtual reality experience. Deploying high-frequency ranges can provide maximum pipe size over alternatives but can come with highest cost-per-square-foot given spectrum range and propagation.

"Spectrum is the essential lifeblood of wireless connectivity—a vital component that makes a network wireless."

“Goldilocks” or “just right” spectrum resides in the mid-band frequencies with prime areas between 1.9GHz to 6GHz that offer a strong balance of propagation, penetration, coverage and capacity with respect to associated ownership needs and restrictions.

It’s important to balance coverage and capacity when aligning frequency ranges needed for your site and user types. It’s also important to understand how different frequencies are aggregated to create larger pipes. While spectrum bands can be utilized without the capital overhead of traditional models, it can come as a cost relating to available or predictable capacity or impact jitter. For example, unlicensed or shared spectrum can be fused with licensed spectrum to increase pipe sizes across a device and radio. Things become more challenging when carrier aggregation is utilized by operators. This can affect an enterprise’s total available capacity when utilizing unlicensed or shared bands.

After comprehending the role frequency plays in coverage and capacity, the next consideration is spectrum access. Spectrum is available in three categories: unlicensed, shared and licensed. These categories add another layer of consideration when determining an optimal solution. The following breaks down key details of each category.

Unlicensed Spectrum

Unlicensed spectrum is free to use and available to everyone. Wi-Fi is a prime example of unlicensed spectrum. Anyone has the ability to purchase Wi-Fi equipment and set up a wireless network. Wi-Fi 5 (adopted in 2014 and supplemented by Wi-Fi 6 in 2019) operates at 2.4GHz and 5GHz, both considered mid-band frequencies capable of “good” coverage and capacity. However, because the spectrum is free to use, regulations stipulate the total power output emitted by a Wi-Fi access point. Aside from some limited exceptions, typical output power is less than 1 watt.

Also, because this is unlicensed spectrum, other devices utilizing similar frequencies and networks can lead to interference by consuming the available “lifeblood.” Unlicensed spectrum requires an asynchronous control plane which creates a “listen first, talk second” mentality. This instability impacts both network reliability (jitter), which can hamper performance capabilities. Reducing the moving speed of a robot for safety reasons is just one example of network instability.

Large areas or sites that might feature interference or materials that limit mid-band transmission may require a larger number of Wi-Fi access points to provide adequate coverage for wireless devices and to mitigate jitter. Here, instead of licensed spectrum, a customer would need to deploy an unlicensed solution to achieve desired performance levels, which can lead to higher overall capital costs.

Unlicensed spectrum can play a valuable role in enterprise-grade solutions for manufacturers, healthcare, retail, entertainment venues and other market segments. With the move to Wi-Fi 6, unlicensed spectrum will co-exist and may converge at a common aggregation point to complement a range of designs featuring hybrid spectrum solutions blended between Wi-Fi and cellular.

Shared Spectrum – Citizens Broadband Radio Spectrum (CBRS)

The second spectrum type is “shared spectrum,” or as it’s better known today in the US, Citizen Broadband Radio Service (CBRS) spectrum. CBRS operates at 3.5 GHz, which is considered a “mid-band” frequency delivering solid coverage as well as capacity. The use of CBRS is managed through the Spectrum Access System (SAS) and allows for higher power output than unlicensed spectrum but still features lower power than licensed. From a control plane synchronization perspective, it does have a “listen first” model, but can maintain session status similar to licensed spectrum once established. Thus, after initial connection, jitter should look similar to a licensed spectrum model.

Before the arrival of CBRS, deploying large-scale cellular networks was not only technically complex but also capital-intensive. CBRS allows entities to deploy private cellular networks that utilize either General Authorized Access (GAA) or Priority Access License (PAL) spectrum. GAA is technically free to use with a Spectrum Access System (SAS) subscription while a PAL will need the SAS plus a purchase of a license during an FCC auction. These developments allow both public and private sectors to include added benefits of private cellular into their local network designs without the typically high capital overhead of licensed spectrum.

Additionally, GAA can access the full spectrum portfolio (up to 150MHz) assuming it isn't in use by a Priority Access License (PAL) holder or an incumbent (i.e., the Navy, a serious consideration near a coastline or large body of water). In these areas, access is managed by the Spectrum Access Service (SAS) and the Environmental Sensing Capability (ESC) as a governor for the incumbents. If an incumbent owner's vessel were to arrive at a port, large swaths of spectrum would be impacted. If the incumbent happens to be leveraging the spectrum in combination with a PAL user, then the GAA network performance sharing can be impacted significantly, and the incumbent will get priority.

Using GAA, any entity with the right network construct and SAS connection can create a private cellular network while following hierarchy from incumbent to PAL. In a rural location outside of the coastal areas, GAA spectrum could be a great solution. But as entities get closer to urban locations, the use of this spectrum will be become riskier from a predictable capacity perspective especially in a coastal city. However, all areas can expect GAA (and PALs) may be used by operators as carrier aggregation as previously described.

In total, there is 150MHZ available per area, about half of it is geared to PAL owners with the remaining used by GAA entities while abiding by the SAS "traffic cop." As customers consider CBRS, they should know the PAL owners, their amount of spectrum blocks as well as the amount in MHz that can be enabled on a particular radio (available total capacity and speed). As an example, a radio may aggregate four 10MHz channels with a total capacity of 40MHz. Additionally, some radios may have a limitation to require contiguous channels to aggregate, and this may not be readily available in a given area.

In the early days of CBRS, capacity won't be as challenged. But, over time as more entities come online, it could become more constrained. With respect to these facts, most entities thinking about private cellular are looking at a multi-year build, making it important to ensure experiences have scale over time.

Licensed Spectrum – Owned/ Carrier-Grade

If spectrum were gemstones, licensed spectrum is a sparkling diamond. It affords owners the most control and usability over other spectrum options, but that also comes with a heftier price tag. Cellular network operators own spectrum licenses that span a range of frequencies, from low-band (sub 1 GHz) to high-band (up to 40 GHz). This allows operators to mix and match frequency types to provide coverage as well as capacity.

> *"A hybrid spectrum strategy—featuring licensed, shared and unlicensed spectrum—can address nearly every challenge from range, penetration and propagation."*

Licensed spectrum carries a synchronous control plane and does not have to "listen first" in any capacity as the entity owns and controls the use of this medium. There's another benefit for licensed spectrum holders as well. Associated radios can operate at higher wattage outputs than those using shared and unlicensed spectrum. Which, by its physical nature, gives licensed spectrum maximum range and distance for each level of band (low, medium and high).

With fewer towers and lower fiber and required infrastructure, a licensed spectrum build can offer an enterprise lower infrastructure costs. In a typical deployment, licensed spectrum power output is typically 1 or 2 watts indoors and can be up to hundreds of watts outdoors. By contrast, shared spectrum is limited to 1 watt indoors and 5 outdoors. Wi-Fi power is usually less than a single watt indoors, and lower power affects distance, penetration and costs.

To gain access to this type of spectrum, a customer needs to "carve out" of an existing holder's spectrum or purchase clean unowned spectrum from a regulatory body (like the FCC in the US).

Spectrum--how to choose?

A hybrid spectrum strategy—featuring licensed, shared and unlicensed spectrum—can address nearly every challenge from range, penetration and propagation. In fact, if a hybrid solution doesn't fit the characteristics when all bands and types are available, customers should consider revising the use-case experience or pursue a wired connection.

A blended approach using multiple spectrum types can begin with a strong, wide-reaching low-band coverage layer that handles a wide variety of fundamental mobility use cases like push-to-talk. Then, add a layer of mid-band spectrum for good coverage with added capacity to allow for capabilities including mobile video collaboration. Finally, strategically blend very high frequency (mmWave) bands that provide ultra-fast throughput and capacity for fixed use cases and experiences including virtual reality. In a perfect world, this full spectrum mix is ideal. But this strategy should keep in mind elements that ensure alignment with devices, radios, and, of course, finances.

Access Network

The radio access network (RAN) is the radio infrastructure that enables devices to receive and transmit information and ultimately connect to the core network. Cellular radio access networks typically consist of antennas with radio heads propagating the wireless signal that connects to a base station and processes the signals, which then converts everything to packets of usable information.

Designing and building the RAN is a critical task. Design variations can greatly impact the infrastructure required to cover the area, which ultimately influences costs and performance. Real-world deployments may have use cases indoors, outdoors, across physical borders or combinations of the three. It's no easy feat to design a solution that provides pervasive campus-wide connectivity blending multiple end points to work seamlessly throughout the campus and off. There are a number of different approaches in these situations.

One example leverages higher power radios for outdoor antennas pointed into the campus to counter similar-frequency towers in proximity. This also provides partial to complete indoor coverage that reduces cost-per-square-foot in fewer fiber and radio units. In this approach, customers should pay attention to building type, other users of a certain spectrum, and, of course, use case experience. Purpose-built radios create an umbrella of RF that can block external usability.

At this point, a customer can add radios for targets not accounted for in the "outside pointed in" build. Whether sub-6 or mmWave, radios installed indoors in this capacity can be used to fill in the gaps. Optimization through physical changes (antenna direction, height, power, etc.) and network parameters can provide seamless user-traffic transition between indoor and outdoor systems that ensures the best user experience. Again, there are many possible approaches. Overall, after given the coverage capabilities of high-power radios especially in the licensed spectrum space, it may make sense to think about building from the outside in to optimize build time and costs.

Transport Layer

The transport layer (i.e., backhaul, fronthaul or even mid-haul) is how data returns to the network core, the compute location or where data is stored and analyzed. Although component solutions can appear simple, transport technology, location, terrain and availability may complicate the solution design process.

Transport options include layer 2 fiber, layer 3 OTT connectivity, cable, microwave, satellite, MPLS, Dedicated Internet, new satellite-based options like Starlink as well as other new concepts such as Integrated Access Backhaul (IAB).

Depending on the critical nature of the deployment or specific business requirements, these options can quickly inform the starting topology. When weighing transport options against KPIs, it's important to note these requirements may alter cost, delivery timetable and what is possible at the location.

Providers can be hindered by limited specifications, a lack of flexibility or by a transport's footprint. Although this landscape is evolving, this is an area where the experience of carriers and broadband providers are leveraged to deploy licensed or unlicensed spectrum in a local, single cloud or hub-and-spoke design.

5G Changes to the Core

Along with spectrum availability changes, the virtualization of the packet core is creating opportunities for operators and enterprises alike. The packet core is the cellular "brain" that powers authentication, handoff and reliability interactions.

As we transition from the evolved packet core (LTE) to the next-gen core (5G), these opportunities will increase. However, similar to previously described spectrum variants, different core element configurations may affect different areas of enablement, performance and scale use cases.

In the past, 4G core providers typically hosted network cores in centrally managed datacenters geo-located throughout their home country or around the globe. When a device would authenticate, all traffic requiring authentication and data would flow from the device through the RAN and transport (or lack thereof) into the nearest core location. Then, that traffic would be handed off to another network or private wide-area network (WAN) and finally to a customer application in their datacenter or the cloud.

This method, while highly secure, reliable and globally scalable, didn't account for enterprise-level control, latency, jitter and reliability requirements within enterprise campus environments. Plus, new technologies including augmented reality (AR), virtual reality (VR), extended reality, remote robotics, autonomous vehicles, drones and many other forward-leaning ideas require a network infrastructure that accommodates greater performance and reliability needs.

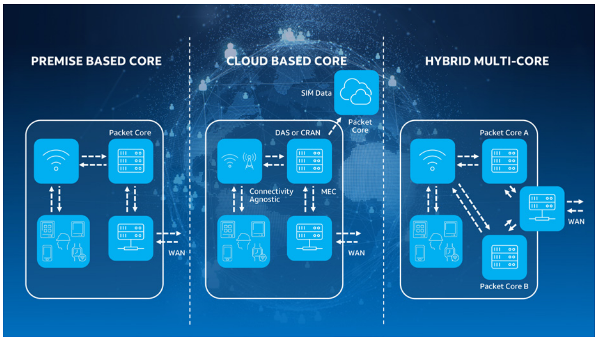

5G has completely reimagined core architecture and its functionality. Key transformations include network functions virtualization (NFV) as well as the ability to separate the control plane from the user plane. This allows providers to bring the core and its functions closer to the edge, thus enabling a host of new service-based architectures including, premise-based, cloud-based, distributed, hybrid or a combination of all three.

Premise-Based Core

The full premise-based cellular core model offers the greatest control and security from the outside world. Configured with the cellular core premise-based, it allows a customer to create an "island" of cellular.

This highest level of cellular network control comes with a comparably hefty price tag especially where multiple sites are concerned. However, the "on an island" approach is key in rural locations where backhaul is limited, difficult or expensive. Other considerations including spectrum ownership also come into play, which can affect cost and capacity. Adding more spectrum may also require government or other-owner support and regulation along with an evaluation of devices.

This island configuration offers great security and control. It also means the amount of redundancy will rely solely on the local engineering and build. The customer owns the network using private Subscriber Information Module (SIM) or Public Land Mobile Network (PLMN), so traffic flows through a highly secure "pipeline" they fully control. Additionally, customers decide how network traffic flows, including adding prioritization and quality of service. Finally, private networks can scale to a great number of users through virtualization. This also may bring higher costs and varies with the spectrum amount required across their user groups.

A premise-based private core is not without its limitations. Additional users joining the network over time can denigrate performance especially as the finite resource of spectrum gets consumed. Scaling in general is a direct function of spectrum holdings and the boundaries of the available core-compute infrastructure. Device availability can be limited as well depending on the frequency range selected. Plus, private networks mandate owners must understand not only how to manage the normal workings of new server and radio infrastructure, but also comprehend cellular differences over alternatives to ensure proper support, optimal experience and scale over time. Customers must also manage the entire Radio Access Network (RAN) node and perform maintenance infrastructure management functions that will most likely be new endeavors for most organizations.

Finally, because these are typically controlled environments, off-the-shelf device availability may be limited. Evaluating a premise-based approach should consider devices as well as other parts of the six-way match that scale in context of spectrum acquired or to be utilized through a shared construct.

Cloud-Based Core – "Hub & Spoke"

A hub-and-spoke model features a centralized packet core that might also include a customer working site or datacenter. The spoke ends are local private cellular LAN environments that may also require mobility of assets between these ends or multiple sites.

Think of a multi-site manufacturer building a widget and moving these widgets to a nearby warehouse. Both sites would need a high-performing LAN for nearly all use cases. With core software costs now shared between sites, a distributed model can help deliver scale. However, to get to an "island" level of high-performing KPIs per site, redundancy, care of backhaul and careful planning in terms of fiber mileage among other things are required. Additionally, physical distances will impact latency requirements given speed-of-light limitations. Core user planes can be distributed, localizing data created at the

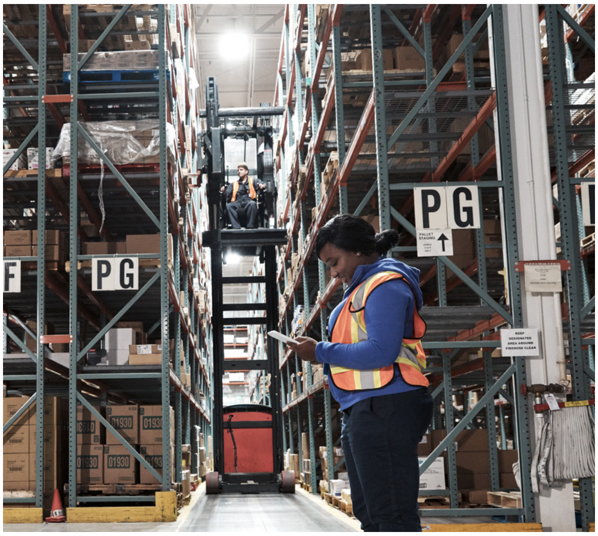

end of each spoke allowing for central control mechanics while allowing optimal performance at the end of the spoke.

When paired with a mobile provider, this model can offer an opportunity for greater spectrum scale without the capital investments needed to own the full range of spectrum from low- to high-band. A customer may inherit not only cellular capabilities, but also the pre-existing scale of an operator while garnering localized routing with other control opportunities, although perhaps not as rich as the "island" approach.

Hub-and-spoke models can also be accomplished with customer-owned or shared spectrum that allows more network control but more limited spectrum capacity due to cost. In considering this model, an enterprise needs to account for the interconnectivity between hub-and-end-of-spoke (transport) and access. These considerations impact cost and performance.

Thanks to new developments in 5G 3GPP standards and the ability to separate the control and user planes, mobile operators can now leverage licensed spectrum to allow both enterprise and government entities to deploy hub-and-spoke designs. These edge-based design solutions deploy licensed spectrum in local area networks indoors and outdoors. Private, highly secure and flexible, carrier-grade spectrum offers high-value propagation and capacity with cost structures atypical of historic operator models. With advanced control, security and data privacy keeping cellular data within a local area network, edge (carrier hub and spoke) builds allow for enterprises to harness the power of a full range of licensed spectrum networks (LTE, 5G sub-6 and 5G mmWave) as well as anticipate tomorrow's network technologies to deliver newfound levels of intelligence, control, reliability, security and speed.

The Hybrid Model

Call this the "Field of Dreams" model: If you build it, they will connect. Hybrid offers a "greatest hits" package for customers that may need to support portable, nomadic, private, public, premise-based and cloud-based architecture (yes, it's already happening). A hybrid model could also be global, redundant, core- and connectivity-agnostic and include a shared Radio Access Network (RAN) or spectrum types that can be leveraged across multiple 5G cores with orchestration. It's worth noting this version of "shared" isn't the same context as previously detailed regarding spectrum.

While we've alluded to top-level constructs, there are several variations of the "island" and "hub and spoke" approaches that accommodate use cases in the six-way match. Even then, it may not prove to be enough. In those cases, ways are emerging to share RAN or spectrum, which is the finite resource among any wireless network.

For these scenarios there may be one-to-many or many-to-many options for core and RAN/spectrum. This accounts for gaps in the prior models but will carry the heaviest cost and complexity. While this could be the day-one model, in many cases this would be a capstone architecture after the individual parts are exhausted based on anything from control to scale.

With complete network control, customers can also take advantage of any combination of spectrum from shared to privately owned licensed within their local area cellular network.

Private PLMN or SIM roaming is yet another compelling concept that could be included to create private cellular network synergy. Imagine a company with a private network that also allows for roaming in and out of a carrier's network when they leave a coverage area. The customer's private PLMN would then "roam" onto a carrier's network and ride a private mobile connection back to their own data center, NOC or cloud-hosted infrastructure.

The list of innovations in the hybrid space is evolving as site types, use cases and business drivers continue to evolve. The hybrid model will provide for unique solutions where cost is not the primary consideration. Flexibility, control and security as well as unique business requirements will inspire building inventive hybrid networks. Interestingly, many of these techniques are already in use by operators who are finding ways to push those environments into customer domains for maximum efficiency and scale.

Last, not least—Content, Application & Solution

The content, applications and solutions make up the last component of the six-way match. The application will ultimately define such network requirements as latency, jitter, downlink capacity requirements, uplink requirements, privacy, security, scalability, reliability and, of course, cost.

One key application architecture trend is computing at the edge to improve performance. As organizations consider this approach, it becomes equally important to also evaluate network traffic routing. Private cellular networks (licensed and shared) now offer an edge approach for the packet flow so data can also be acted upon at the edge. Tightly coupling cellular networking and computing at the edge is an architecture for organizations to consider driving higher levels of performance and improve the end-user experience while possibly reducing the space requirements of the physical infrastructure. In this paradigm shift, mini-cellular networks with associated compute and application nodes are built to allow for greater performance or data isolation. But, again, this must be weighed against experience, performance and cost.

As we close out the six-way match section, it's important to note that no single use case will validate a shift to "insert new wireless LAN" network architecture. In many cases, the shift is made by assuming most users and things can benefit from simple raw connectivity. Another reminder that data typing is a highly valuable initial step. Secondarily, it's also often safe to assume a virtualized and spectrum construct allows Wi-Fi to perform similarly to cellular.

The six-way match has offered an effective way to educate and inform our customers as well as guide our internal product roadmaps and designs. However, before any options are finalized, security and privacy must be reviewed in parallel to each piece part from device to application.

"Choosing a partner that can help navigate all the different choices with experience and depth across the entire private cellular spectrum is critical."

Security & Privacy Matters

As noted, while not highlighted specifically in the six-way match, security and privacy are an integral part of each point. Integrating data security and privacy into networks complements the six-way match by ensuring a system-level approach is applied to security.

According to the 3GPP standard, cellular networks add new inherent security layers. Air interface encryption up to 256-bit encryption with 5G-based private cellular networks, GTP tunneling protocols for every session and now the ability to keep user-plane traffic on-premises offer higher levels of privacy.

One key area to consider along with knowledge of the inherent security capabilities is data-handoff between the cellular user plane or packet core and the existing LAN or application server. To extend this data to the customer's private enterprise network, connectivity is established through building an APN (Access Point Name) represented by the packet core features. Once an APN is established with the customer's IPs, data can be handed off similarly to methods used by carrier networks when going to the internet or a private network like MPLS. Other handoff models can be executed with VPN (IPsec or GRE tunnel), allowing customers to more securely interwork with their internal services with end-to-end packet encryption that becomes highly important with devices that flow in and out of a campus using multiple networks.

Private network applications—hosted in cloud-based or standalone premise-based servers—are further protected by defining network or premise-based firewall rules at the newly created internal DMZ providing access that can replicate existing security protocols.

Overall, once a good understanding of inherent security is in place, we recommend viewing traffic flow from cellular devices to and from the existing network into an enterprise LAN or app server. This method helps security teams find their policy balance at a familiar starting point (aligned with existing company policy) and expand based on user context, risk and experience.

Greater Control

The last three years, we created a hypothesis of using cellular more efficiently by employing new virtualization and spectrum options. We realized high-performing wireless connectivity solutions—to integrate people, sensors, machines and more—would continue to grow as enterprises moved to Industry 4.0 and required more from their environments. Our solution? Create a suite of cellular solutions with basis and growth around the six-way match. Our methodology can help you navigate your campus networking needs from a single site to multi-national environments.

Staying true to the six-way match while using existing macro systems and in-building topologies, helped us create three private solutions that offer varying control levels. These constructs can add value and allow customers to deploy solutions for today with resilience for the future.

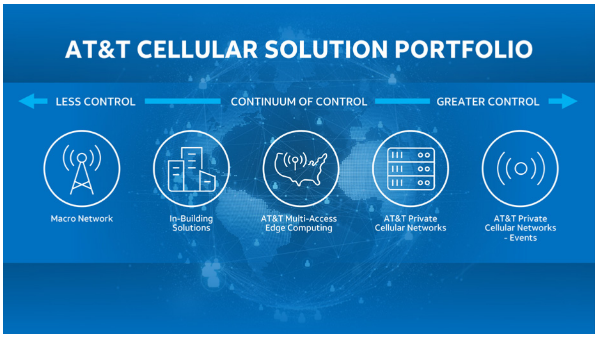

AT&T: The Perfect Partner For The Six-Way Match.

Given the many pieces of a private cellular solution, understanding each six-way match component is critical to developing optimal solutions. In this new and fast-growing area, a brand's experience and expertise should impact the decision-making process. As your organization thinks through these considerations, ask: What are all the pieces we need to bring together to meet the needs of our business? Can we support and scale the solution? What are the hidden costs of recruiting and retaining the right skill sets for these next-generation wireless networks? What are the opportunity costs across all processes?

Choosing a partner that can help navigate all the different choices with experience and depth across the entire private cellular spectrum is critical. When access, speed, security, control and connectivity count, experience counts the most.

With our proven experience, AT&T is uniquely qualified to deploy and manage a wide range of wireless network types. Few can match our hand-on expertise in every aspect of the six-way match—from device to spectrum to core to fiber and globally. An ability to provide network options including private cellular at scale in multiple ways, affords enterprises the ability to drive optimal connectivity across use cases and site types.