WHITE PAPER

CBRS, 5G and Wi-Fi: Radio Access Network Convergence in the Enterprise

aruba
a Hewlett Packard
Enterprise company

TABLE OF CONTENTS

EXECUTIVE SUMMARY 3

ARUBA’S VISION OF THE MULTI-RAN ENTERPRISE 3

PRIVATE CELLULAR DATA & VOICE 4

ENHANCING IN-BUILDING PUBLIC CELLULAR NETWORK COVERAGE 5

CURRENT STATUS OF IN-BUILDING CELLULAR RAN TECHNOLOGIES 6

CITIZEN’S BROADBAND RADIO SERVICE (CBRS) 12

WI-FI CALLING & ARUBA AIR PASS 14

CELLULAR DATA PATH CONVERGENCE WITH ENTERPRISE NETWORKS 16

5G IN THE ENTEPRRISE 20

CONCLUSION 22

EXECUTIVE SUMMARY

Enterprise networks employ a diverse set of physical network types for access, distribution and backhaul. The larger the enterprise, the more physical network types that are in use. From an enterprise perspective, cellular radio technologies including privately-owned LTE – of which the kind enabled via the Citizens Broadband Radio Service (CBRS) spectrum in the United States is one example – and eventually privately-owned 5G can be best understood as essentially special-purpose access networks with particular technical, network topology and cost characteristics. They are useful additions to the enterprise network architect's toolbox to solve specific business problems in addition to or in parallel with more traditional access networks like Wi-Fi or Ethernet. Wi-Fi, and CBRS-based LTE/5G wireless – as well as new IoT radios such as Zigbee – are all radio access networks (RANs) and they will be increasingly combined by enterprises in the coming years. In other words, enterprises will operate in multi-RAN environments.

Enterprise cellular RANs are built with small cells, like Wi-Fi access points (APs). They can provide data service, voice service or both. These services may be provided on a "private" basis where the enterprise itself owns the subscriber identity module (SIM) credential, or on a "public" basis where roaming privileges are extended to subscribers of well-known mobile network operators (MNOs). Depending on the use cases to be served, cellular RANs can be deployed independently of Wi-Fi, or may be co-deployed with coverage engineered to match the Wi-Fi footprint to support client devices with both types of radios.

Over the next few years, the most promising enterprise uses of cellular RAN technology are for private voice and data applications. This includes mobile point-of-sale, internet-of-things (IoT), push-to-talk voice and warehouse automation. In these scenarios, the enterprise owns the entire end-to-end system.

Providing public voice roaming services between macro cellular networks and privately-owned enterprise RANs on a neutral host basis is challenging as of this writing. There are business constraints – primarily the establishment of roaming agreements – as well as technical challenges ranging from handover signaling to establishing the necessary secure connections to operator core networks. Solutions are being worked on collaboratively within the CBRS Alliance on the one hand, and within the Wi-Fi Alliance on the other. We will explore these challenges in some detail in this paper, and the potential timeframe for solutions.

This white paper provides an overview of the state-of-the-art of enterprise RAN technologies including both cellular and non-cellular variants. It explains how they may be combined from a control plane and data path perspective in an enterprise network architecture. Our objective is to arm customer IT decision-makers and architects with the basic conceptual knowledge, terminology, architectures and deployment options to facilitate strategic planning.

ARUBA'S VISION OF THE MULTI-RAN ENTERPRISE

Enterprises have long had diverse connectivity requirements at the edge. Since at least 1990, Ethernet has been the dominant wired access layer in the enterprise, with Wi-Fi supplanting it for most indoor mobile broadband use cases since the rise of BYOD and the introduction of the iPhone in 2007. Recently, new special-purpose wireless access networks have begun to be deployed in the enterprise. Bluetooth Low Energy (BLE) tags in unlicensed 2.4 GHz spectrum are being used for an array of asset tracking and wayfinding applications. Zigbee, Thread, ISA100 and WirelessHART – all based on the 802.15.4 standard using unlicensed 2.4 GHz or 900 MHz spectrum – have become de facto in-building IoT access networks with Zigbee alone having over half a billion installed units worldwide.[1] And Wi-Gig is a new ultra-high speed broadband technology operating in the unlicensed millimeter wave bands at 60 GHz. To borrow the cellular term, each of these are separate enterprise radio access networks (or "enterprise RANs").

Today's dominant enterprise network architecture weaves these elements together in specific ways to serve the four principal enterprise workload geographies depicted in Figure 1: campus, regional office, branch and small office/home office (SOHO). Some enterprises with significant populations of vehicular and travelling users have a fifth geography, which uses a wireless WAN (WWAN) using 3G or 4G modems from one or more mobile operators.

The enterprise wired and wireless "edge" is shown at the bottom of Figure 1. These are all the different devices that are served by the access layer of the enterprise network.

[1] https://zigbeealliance.org/news_and_articles/zigbee-leads-the-wireless-mesh-sensor-network-market/

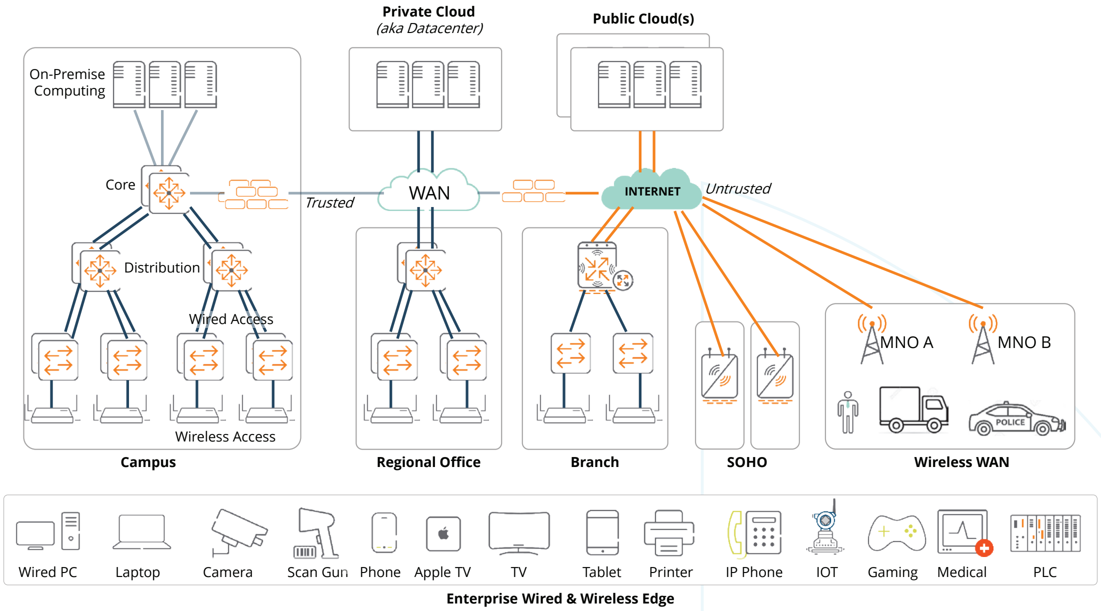

Figure 1: Classic Enterprise Architecture Serving Four Primary Fixed Workload Geographies Plus Mobile WWAN Users

Devices such as laptops, phones and printers may exist in every geography. Certain devices such as point-of-sale scanning guns and autonomous guided vehicles (AGVs) or medical equipment or IoT may only be found in certain parts of the network. It is increasingly common to find some classes of devices that include support for two or more RANs in use at the same time. A smartphone with LTE, Wi-Fi and Bluetooth is a familiar example, but headless IoT devices such as medical devices or portable scan guns increasingly come equipped with multi-RAN features.

Against this backdrop, private 4G and 5G can be best understood as special-purpose access networks with specific performance and cost characteristics to serve particular application use-cases or groups of end-user devices. These additional forms of enterprise RANs are useful additions to the network architect's toolbox to solve specific business problems.

In short, Aruba sees the enterprise edge as:

- Composed of multiple, overlapping enterprise RANs using different radio technologies to serve particular, differentiated device types and use cases
- Leveraging unlicensed, shared spectrum (e.g. CBRS) and licensed spectrum as appropriate
- Including targeted use of wired Ethernet to directly connect certain classes of end device as well as the radio heads
- Segmented into enterprise trust domains, with transit between each RAN and upstream destinations controlled by an automated policy framework
- Leveraging both enterprise authentication stores (e.g. Active Directory, ClearPass) and external third-party identity providers (e.g. AT&T, Verizon, Facebook, Apple) as appropriate to the use case
- Unified under a common management framework, network data plane and security policies

Aruba believes that there is no one-size-fits-all answer to access-layer connectivity at the edge. We are committed to a holistic approach that integrates cellular and non-cellular technologies over time.

PRIVATE CELLULAR DATA & VOICE

The ability to deploy privately-owned cellular RANs in service of enterprise needs is one of the most exciting developments in recent years. This has been made possible by key architectural enhancements by 3GPP[2] and the European Telecommunications Standards Institute (ETSI) to permit integration with the enterprise data path instead of having to route traffic to an MNO core network.

[2] The Third Generation Partnership Project (3GPP) is a standards body for the cellular industry, similar to the role that the Institute of Electrical and Electronics Engineers (IEEE) plays for the enterprise network industry. Ethernet and Wi-Fi are IEEE standards, whereas LTE and 5G are 3GPP standards. 3GPP cellular standards are known as "Releases" and are numbered. For example, the first 4G standard was initially introduced in 3GPP Release 8 in 2008. The 5G standards were initially introduced in Release 15 and are being extended in Releases 16 and 17 which are due in mid-2020 and mid-2022, respectively.

The ability to route IP traffic between cellular client devices and enterprise networks opens a range of applications across numerous verticals, such as:

- Retailers can operate mobile point-of-sale terminals and inventory scanners, connect building IoT systems, feed ruggedized tablets on forklifts, and power robotic warehouse systems
- Manufacturers can use wirelessly enabled power tools to record every aspect of a product's creation as it moves down the line, and apply machine vision systems for automated quality inspection
- Public venues can perform ticket scanning, enable push-to-talk (PTT) voice communication between staff members, and sports teams can provide secure sideline data terminals for real time decision making
- Hospitals can deliver latency-sensitive medical telemetry to nursing stations and electronic medical record servers, and provide PTT voice communication for clinical staff, and enhance in-building cellular service for patients, families, and staff

Wi-Fi has been successfully addressing virtually all these same use cases for over 20 years. However, some Aruba customers would like to augment their WLANs with cellular technology for specific reasons, such as:

- **Reserve 100% of Wi-Fi spectrum for guests, fans or another user community.** For example, a theme park or stadium may wish to offload all its back-of-house networking to cellular to maximize capacity for fans.
- **Congested or dirty spectrum.** Some enterprises operate in very dirty or unpredictable radio environments. An electronics retailer with a store full of devices generating unwanted traffic may have trouble with point of sale equipment. Airport gate areas experience huge Wi-Fi usage immediately prior to boarding which can interfere with airline devices.
- **Outdoor wide-area coverage.** Due to significantly higher power levels permitted by the FCC, a single outdoor small cell can cover a large area – up to several square kilometers – by itself. This provides a cost-effective solution for industrial sites, airports, oil fields or other large facilities to deliver wide-area data and voice services.
- **More spectrum.** As enterprises go all-wireless, some are finding that even with clean spectrum they simply do not have enough to meet all their requirements. For example, some customers wish to backhaul HD security camera video without tying up Wi-Fi channels.

The multi-RAN enterprise does not have to choose. It can deploy all the RANs that are appropriate to its application needs and radio environment. In some cases, cellular RANs will be purpose-deployed in a portion of a customer's floor plan, such as on a manufacturing line, a football sideline, or at an airport gate. In other cases, the cellular and Wi-Fi RAN footprints will be engineered to match, such as in a retail store where both shoppers and associates can travel the entire facility. The enterprise network architect can adjust the deployment to control cost in line with the use case requirements.

ENHANCING IN-BUILDING PUBLIC CELLULAR NETWORK COVERAGE

Today, every organization requires reliable cellular connectivity from major mobile network operators in and around their facilities. This requirement includes seamless hand-in / hand-out of voice calls that are in progress, as well as voice, data or messaging sessions initiated inside a building. Employees, customers, vendors and guests depend on SIM-enabled mobile devices more than ever in today's world. The coming transition to 5G is expected to unlock new use cases that will only magnify this need.

Unfortunately, commercial buildings have widely varying levels of coverage from MNO macro networks. The two most important reasons for this are geography and building construction. Commercial buildings located in areas with significant radio-absorbing ground clutter such as medium rise buildings or tall trees, or in any area where cellular macro coverage may be sparse such as suburbs often exhibit this problem. The aggressive worldwide adoption of thermally efficient building codes has driven widespread use of low-emissivity ('low-e') coated glass, radiant barrier insulation, thicker walls and other techniques that pose a fundamental challenge to the penetration of the wireless signals in the frequency bands currently used by 4G LTE systems.

This problem is expected to intensify during the 5G transition. That is because the "core" 5G spectrum is in the upper 3 GHz range. By contrast, most existing 4G networks operate at or below 2.1 GHz. Enterprise WLAN engineers are well acquainted with the propagation differences between the 2.4 GHz and 5 GHz bands used by Wi-Fi today. Compared to the 2.4 GHz band, 5 GHz signals experience a minimum factor of four times (6 dB) more propagation loss. Higher frequency signals are also more severely attenuated by common building materials. 5G networks using mid-band spectrum will suffer the same problem as compared with 4G systems. Unlike Wi-Fi which only must deal with in-building attenuation, public cellular networks historically need to penetrate the building shell which results in vastly higher attenuation for the reasons mentioned above.

Recent developments in small cells including the new CBRS band in the United States offer a possible solution by moving the cellular radio connection indoors. As stated in the introduction, there are both technical and business challenges that must be solved to offer a truly neutral host public cellular experience on privately-owned small cells. We will look at this in detail later in this paper.

Another alternative that shows real promise is Wi-Fi Calling combined with Aruba's Air Pass service. Air Pass allows SIM-enabled smartphones to detect and automatically connect to the existing guest Wi-Fi network using the subscriber identify module (SIM) credentials provided by a mobile network operator (MNO), making transparent Wi-Fi Calling and messaging services available. We will investigate both solutions later. It should be noted that both CBRS and Wi-Fi Calling with Air Pass can be combined to reach even more devices. On the one hand, CBRS band support will take some years to become widespread. On the other, Wi-Fi Calling must be manually enabled due to e911 rules and so adoption is not universal. Deploying both Air Pass and CBRS together can maximize the service quality.

The bottom line is that despite many advances in enterprise and campus communications in recent years, and the extensive build-out of outdoor cellular networks, many enterprise customers continue to suffer from poor in-building cellular coverage. They seek a simple, inexpensive, universal solution. Aruba recognizes that there is no one-size-fits-all answer to this problem, and no "silver bullet." We have been closely studying developments in cellular convergence, and we are investing in a roadmap that offers customers a range of solutions.

CURRENT STATUS OF IN-BUILDING CELLULAR TECHNOLOGIES

Historically, the established solution for indoor cellular coverage is to install a DAS (Distributed Antenna System). DAS is well-understood and effective, but very expensive, and most enterprises find the price point out of reach. DAS owners have been seeking lower-cost alternatives for years with little success. DAS is only cost-effective for facilities with hundreds of thousands or millions of square feet so we will not consider it further here.

A significant milestone for in-building cellular occurred in late 2019 when the Federal Communications Commission (FCC) in the United States opened a new band at 3.5 GHz for CBRS, based on a novel 'three-tier' shared spectrum model. CBRS is intended to allow enterprises and organizations to set up private base stations using small cell technology. The band is particularly well suited to the multi-operator requirements, as it does not encroach on any existing operator spectrum and can be deployed as a 'neutral host' service. CBRS products began shipping in 2019, and many new form factors of devices are expected soon. We will look at CBRS in depth shortly.

Small Cells

On-premise cellular RANs are built using small cells. A small cell plays a similar role in a cellular architecture that Wi-Fi access points do in a traditional network. Small cells come in similar form factors to Wi-Fi APs, use Power over Ethernet (PoE) for data and electricity, and are intended to mount to a ceiling grid, wall or roof truss. They come in both indoor and hardened outdoor varieties.

At its simplest, a small cell is a miniaturized base station. It has a low-power radio with a range of perhaps 100 meters and reduced capacity in terms of served users, but in other ways the radio functions like a macro base station.

Most small cells have two radio slots. Just like a Wi-Fi AP has separate 2.4 GHz and 5 GHz radios, a small cell will usually have two different frequency bands. But unlike Wi-Fi which has only two fixed bands that are globally harmonized, small cells have tunable radios that can be adapted to most or all the defined 3GPP bands.[3] As of this writing, there are nearly 60 different 3GPP bands. CBRS is one such band, known as "band 48" (or "B48"). A 3GPP band may be frequency-division duplex (FDD) or time-division duplex (TDD), may use various channel widths from 1.4 MHz up to 20 MHz, and may require a license or may be shared spectrum. We will look at 3GPP frequency bands later in this paper.

Typical deployment densities for 4G LTE and 5G NR small cells operating at or below 3.7 GHz are quoted at roughly one per every 10,000 ft^2 (1,000 m^2) for indoor deployments. This is approximately four times less dense than Wi-Fi systems where each AP covers approximately 2,500 ft^2 (250 m^2) accordingly to long-established manufacturer design best practice.

[3] https://en.wikipedia.org/wiki/LTE_frequency_bands

However, if the goal of the deployment is to provide in-building cellular services with high device density then the total number of licensed low- or mid-band small cells may not be much different than the Wi-Fi AP count. This is because of the need to have *separate layers* of small cells for different cellular operators. Since all four major US operators today require discrete small cell hardware, a four-operator deployment using licensed spectrum would require as many small cells as Wi-Fi APs. This unrealistic cost structure is one of the key drivers behind the push for neutral host spectrum such as CBRS.

As of this writing, small cells require between 23 watts and as much as 50 watts of continuous PoE depending on the manufacturer and number of radio slots.

"Private" vs "Public" Cellular Networks

One of the single most important concepts to understand in evaluating the suitability of a cellular RAN for an enterprise use case is a "private network" versus a "public network". In simple terms, this refers to who issued the SIM cards being used by the devices and whether subscribers of an MNO or only authorized devices of an enterprise can access the network.

The defining characteristic of a private LTE network is that the SIMs are issued by – and managed by – the enterprise. In the future, "private 5G" networks will also share this attribute. Enterprises will be able to build their own LTE/5G networks with SIMs provided by the product vendor with the necessary unique encryption codes on them which can then be inserted into compatible endpoint devices.

Given this definition, it comes as no surprise that a "public LTE" or "public 5G" network uses SIM identities that are issued by a major public cellular network operator such as AT&T, Vodafone, Verizon, NTT Docomo or others.

These apparently simple definitions yield profound differences in network architecture, cost and complexity for cellular RANs in the enterprise. These include:

- Devices cannot use a privately issued SIMs to connect to a public cellular network.
- Devices cannot use publicly issued SIMs to roam onto a private cellular network without extensive technical integration and complex business agreement(s) between the SIM issuer(s) and the network operator.
- Publicly issued SIMs are inherently untrusted from an enterprise network security point of view because they have no corresponding enterprise identity.
- Privately issued SIMs may be trusted by the enterprise since they can be tied to a specific role policy and the root of trust is owned by the enterprise.
- Privately issued SIMs require a compatible 4G Home Subscriber Server (HSS) or 5G User Data Repository (UDR) to be deployed by the enterprise as part of its privately-owned cellular solution.
- Publicly issued SIMs do not require the enterprise to operate an HSS or UDR. However, the core network implementations of each supported MNO must support the necessary 3GPP interfaces for "visited" networks to communicate with "home" networks for authentication and billing purposes.

Dual-SIM smartphones are beginning to become more prevalent, either with two physical SIM slots or a single physical SIM combined with eSIM functionality. These phones can have both a public MNO SIM and private LTE SIM. For example, enterprise-provided devices can be served by private LTE when they are on premise and then seamlessly roam to a public MNO network when they move off site. However, it should be clear that in this case the phone is merely an endpoint device in two different networks and the dual-SIM functionality does not bridge the operator and enterprise trust domains.

As a practical matter, this means that if your primary interest in a privately-owned LTE/5G network is to improve your in-building cellular coverage for public MNOs, this objective cannot be achieved without engaging a third-party neutral host provider that has negotiated roaming agreements with the MNOs you require. In addition to those agreements, the neutral host vendor must also have operational VPN tunnels to and core network integrations with each operator to pass control plane messages.

Small Cell Deployment Scenarios

For purposes of this paper, enterprise small cell deployments can be divided into four basic types as shown in Figure 2. Each of these four architectures might use the same small cell hardware (if the manufacturer supports it) but is configured differently in software and has very different cellular core network designs. Three of the architectures employ a single "layer" of equipment serving a single radio channel. One architecture employs four separate "layers" of equipment – one for each major MNO – to provide service using four discrete radio channels licensed by each participating operator.

Option 1: Private Single-Layer Enterprise Deployment

Falling back on our definition of *private vs. public* cellular RANs, the leftmost model in Figure 2 shows a purely *private*

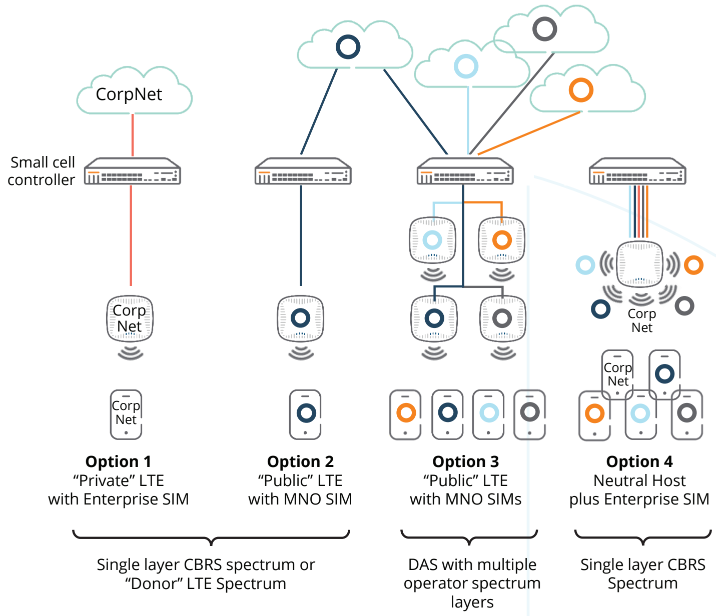
CorpNet
Small cell controller
Corp Net
Corp Net
Corp Net
Corp Net
Option 1
"Private" LTE with Enterprise SIM
Option 2
"Public" LTE with MNO SIM
Option 3
"Public" LTE with MNO SIMs
Option 4
Neutral Host plus Enterprise SIM
Single layer CBRS spectrum or "Donor" LTE Spectrum
DAS with multiple operator spectrum layers
Single layer CBRS Spectrum

Figure 2: Four Principal Small Cell Deployment Architectural Models

single-layer small cell deployment that would serve just a single enterprise. Enterprise SIMs are issued and maintained by the organization, usually with the help of a third-party SIM provisioning company. It is also necessary to deploy a private cellular core network as part of the solution architecture since the SIM identity store is entirely self-contained within the enterprise. This is often provided as a cloud service, sometimes referred to as "cloud RAN" but can also be deployed purely on-premise if required.

A private cellular RAN requires a dedicated spectrum layer. There are two main approaches here. It could be *donor spectrum* that is subleased from a mobile operator or other holder for a fee. Such spectrum cannot be in use anywhere in the area to avoid interfering with the operator's macro network. In practice, this is extremely difficult to achieve because spectrum is expensive and congested, so operators rarely allow it. Second – and much more promising – is a new shared radio band called the *Citizens Broadband Radio Service* that has become available at 3.5 GHz in the United States. A few other countries have begun commercializing similar models. For example, Germany has allocated an "industrial" band at 3.7 GHz open to certain enterprises and similar models are now available in the UK and Japan.

Standalone private cellular networks with private SIMs cannot support roaming of public SIM devices from mobile operators because they lack the roaming agreements, secure connectivity to MNO core networks, and ability to advertise public operator network identifiers. So, Option 1 does not address the concern of improving the public cellular connectivity experience for mobile phone users in commercial buildings other than dual-SIM devices. But a private, closed CBRS network may provide a vital new tool for network architects to support devices that could deliver improved performance from dedicated spectrum. This is the most promising area for enterprises to pursue in the next few years.

Option 2: Public Single-Layer, Single-Operator Deployment

Option 2 is a public single-layer deployment that generally uses licensed spectrum owned by the MNO. Careful coordination of radio channels and roaming is vital, and typically is handled by the MNO. For example, indoor system transmit power levels must be tuned to ensure proper hand-in and hand-out from the macro network to the small cells and back again. Such systems have the advantage of being managed by a single operator but are of limited utility where a heterogenous mix of subscribers to different MNOs is expected in the coverage area.

Mobile operators are also expected to offer CBRS in an Option 2 configuration for their own subscribers. This may involve augmenting the macro network outdoors with supplemental "infill" coverage in hotspots or could entail providing managed services indoors to enterprises. In this case, the small cells would be managed by the MNO and integrated with their macro network via an S1 (4G) or N2/N3 (5G) interface. On paper this is a compelling idea because the spectrum does not conflict with the macro network and thus avoids some of the challenges of coordinating small cell power levels to avoid causing disruption to the macro network. In practice, most enterprises will find it does not meet their needs. Aruba's customers have made clear to us that they want to solve the cellular coverage problem for *multiple operators* not just a single MNO. While this might work for a small or medium sized business with operations in a single city with a single dominant operator, organizations with operations around the country must support all the major operators.

Option 3: Public Multi-Layer, Multi-Operator Deployment

This brings us to the third deployment option in Figure 2 which shows a *public multi-layer* deployment. This permits multiple operators to provide service, where each MNO uses its own spectrum on dedicated equipment. This is the oldest and most established in-building architecture – traditionally built in the form of a DAS – but is also the most complex and expensive. Because each MNO uses licensed spectrum it purchased at auction or was allocated by a national government, and there are typically at least four major MNOs per region, this model requires a separate layer of small cells for each MNO. Public multi-layer systems cannot be easily deployed outdoors because they use the same spectrum as the macro network and will interfere.

This scenario is identical to a DAS deployment, but today can be constructed with lower cost small cells to serve commercial buildings that do not need the capacity of a DAS. There are a few different options in a 3GPP system to affect an Option 3 solution. The prerequisites are:

- The mobile core software and radios purchased by the enterprise must be certified by each MNO as compatible with their macro network
- A roaming agreement must exist between the enterprise and each MNO to be supported
- A secure VPN tunnel must exist between the enterprise and each MNO to pass control plane traffic to coordinate handoffs with the macro network
- If required by the operator, a secure VPN tunnel must exist to pass data plane traffic to home network (aka "home routing"). This is common in some countries in Europe, the Middle East and Asia

Option 3 is depicted in Figure 2 using a single common radio controller from a single manufacturer which manages all the individual radios, and in turn maintains separate control plane links to the various participating MNOs.

However, different MNOs may not all support the same hardware providers. This deployment mode requires active cooperation of all the key mobile operators in your region/country. For a multi-national enterprise intending to roll out a common architecture worldwide, this can be very challenging. MNOs must first agree to permit roaming by their subscribers onto your system. This requires a complex legal agreement with each operator to give you the right to advertise their network on your equipment. You must also contract with a mobile roaming exchange provider and install secure network interconnects to pass authentication and control signaling between the small cell network and each of the participating MNOs. MNOs typically have specific network quality requirements and service level agreement minimums defined in the contract, and they can choose to discontinue participating if these are not met.

Option 4: Hybrid Single-Layer, Multi-Operator Neutral Host Deployment

This is a novel architecture that is possible in certain countries with certain operators. The purpose remains to provide high quality indoor cellular service for multiple operators. However, in this model all the operators agree to share a new radio channel that does not conflict with existing holdings. Because it doesn't interfere with any macro network, it can work indoors or outdoors. The catch is that such spectrum must exist and be assignable to enterprise users. CBRS in the United States is the best example of this.

In Option 4, a technology called *multi-operator core networking* (MOCN) effectively combines Options 1 and 3 together. This permits advertisement of both a private network identifier and one or more public network identifiers on the same physical small cell. With the opening of the CBRS band, the door has been opened to this model. Mobile operators are currently experimenting with MOCN. Such a network would require both the on-premise private cellular core for private enterprise SIMs from Option 1, as well as all the public operator SIM infrastructure and agreements from Option 3.

	Option 1 Private Single Layer	Option 2 Public Single Layer	Option 3 Public Multi-Layer (Multi Operator)	Option 4 Hybrid Single Layer (Neutral Host)
SIM Provider	Enterprise	MNO	Multiple MNOs	MNOs & Enterprise
PLMNID Advertisement	Enterprise	Single MNO	Multiple MNOs	Multiple MNOs & Ent.
Spectrum used by small cells	CBRS Donor spectrum	CBRS Shared Macro	Shared Macro	CBRS
Indoor/outdoor	Either	CBRS – Either Shared – Indoor	Indoor only	Either
Macro network roaming	No	Yes	Yes	Yes
MNO roaming agreement(s)	No	Required	Required	Required
MNO radio coordination	No	Shared macro only	Yes	No
MNO core network integration	No	Via S1 / X2 (4G); or Via N2 / N3 (5G)	Via S1 / X2 (4G); or Via N2 / N3 (5G)	MOCN Via S6a / S8 / S9 Via N3 / N8 / N11 / N15

Table 1: Summary of Small Cell Deployment Models

A second challenge is the issue of subscriber mobility for hand-in and hand-out from macro networks to on-premise CBRS systems especially for voice calls that are in progress to avoid call dropouts. CBRS spectrum does not conflict with any existing MNO spectrum and so does not require any kind of radio-level coordination. However, the MNO core network must still decide to direct and/or permit a Band 48 (B48) smartphone to roam from one network to another. This is a critical difference between Options 2 and 4. With Option 2 the MNO manages its own small cells, which appear to be part of the MNO's overall network. In Option 4 the infrastructure is privately-owned by the enterprise and is not integrated with any MNO network from a roaming perspective.

The technical standards to facilitate roaming under these conditions exist on paper but are just beginning to be implemented in an enterprise context. The 3GPP and CBRS communities are actively partnering with mobile operators to conduct the necessary R&D. Solving this problem may also require that mobile device vendors create roaming firmware to cause smartphones to detect local privately-owned CBRS networks and permit seamless handoff of calls from a macro RAN to privately-owned small cells.

Enterprises that ultimately want to serve both private and public SIMs can get started today with Option 1 as there is no reason to delay while making sure to engineer the RF design so the wireless network is properly dimensioned to meet macro cellular network coverage minimums. An Option 4 deployment could require more small cells than Option 1 in some use cases.

Summary of Small Cell Deployment Models

In this section, we have broadly categorized on-premise small cell architectures available to enterprise buyers into four different types. Table 1 presents a consolidated summary of certain key attributes of the four deployment models. To be clear, each model has a rich variety of individual flavors. In practice, they are part of a continuum of options of which we have just scratched the surface. But this simplified view is useful to understand the broad options as well as the associated complexity and magnitude of likely cost for each one.

3GPP Bands

Unlike Wi-Fi that is globally harmonized in three frequency ranges, cellular devices are not guaranteed to support every 3GPP band. For public networks, this is not usually a problem for longstanding MNO bands but could well be an issue if you need a comparatively new band like band 48 (B48) for CBRS. Table 2 shows LTE bands in use by the four major U.S. operators, along with some European, Chinese and Indian operators.

New bands are added all the time and may well not be supported in all but the most recent handsets (and even then, support may be limited). For example, in 2019 AT&T added Band 14 (B14) for FirstNet, the nationwide first responder network for which it won the contract with the US Federal Government. Also, in 2019, T-Mobile added Band 71 which includes the 600 MHz spectrum it purchased at an auction that year.

Operator	4G LTE Bands	Main Frequencies
US (AT&T)	2, 4, 5, 12, 14, 17, 29, 30, 66	1900, 1700 abcde, 700 bc
US (Verizon)	2, 4, 5, 13, 66	1900, 1700 f, 700 c
US (T-Mobile)	2, 4, 5, 12, 66, 71	1900, 1700 def, 700 a, 600
US (Sprint)	25, 26, 41	1900 g, 850, 2500
Europe	3, 7, 20	1800, 2600, 800
China, India	40, 41	2300, 2500

Table 2: Example 3GPP Bands Used by Major Mobile Network Operators[4]

If you are in the U.S. and interested in a public or private single-layer deployment using CBRS, then you require support for band 48 in both the small cell and every endpoint device. The first premium smartphones with B48 support shipped from Apple, Google, LG and Samsung in 2019. A rich ecosystem of B48 devices and modules has begun to ramp up, with many recent announcements.

CBRS small cells may include two B48 radios, with each radio tuned to different channels. Using the *same carrier aggregation* technology used for LAA, it is possible to combine the capacity of the two radios together into a single larger pipe.

Operator Core Network Integration

Small cells are not usually connected directly to the operator's core network, rather they have a controller that is used to manage functions like inter-small-cell handovers and frequency selection, while offering a consolidated interface to the core network. This separation is necessary because a standard core architecture cannot easily manage the huge number of small cells deployed inside numerous private facilities, but it introduces some challenges.

Perhaps the greatest challenge is multi-operator support. We have already considered the complexities of an Option 3 deployment with separate equipment and coverage layers. Operators prefer Option 3 to ensure guaranteed access to spectrum they control. But with the emergence of CBRS and the intensifying indoor coverage problem, US operators have been more open to exploring Option 4. The Small Cell Forum lists several architectures that facilitate multi-operator sharing of a single-radio small cell and in particular MOCN[5], but so far just a few operators worldwide have developed the technical means and commercial agreements to allow such sharing. There are several US trials of MOCN under way with tier one operators. As mentioned above, in Option 4 the RAN operates on a separate frequency than each MNO's individual licensed spectrum. This can simplify MOCN deployment since there is no interference coordination with macro RANs.

A second challenge stems from connectivity requirements to backhaul small cells to each operator core, especially if the small cells are privately owned. Operator-owned small cells are typically connected via fiber or point-to-point microwave to the MNO's own core network. But a building or campus full of privately-owned small cells is inherently untrusted from the MNO's perspective. Hence the need for legal agreements and secure VPN tunnels for control plane and sometimes data plane transmissions.

Ownership, visibility, management and control can also be problematic. Since the operator is legally responsible for transmissions in its licensed spectrum, it must maintain some level of control. IT managers must deal with small cell radio units mounted on the ceilings of their buildings, backhauled over their LAN but managed by a remote entity. This can become a greater issue when re-configuring or troubleshooting.

For these reasons, the "neutral host" model is likely to enjoy the greatest success for enterprises with interest in public cellular augmentation. A neutral host provider can likely bundle multiple operators from both a network and business perspective, insulating the enterprise from the legal, technical and operational burdens that may well otherwise be insurmountable. And in the near term, enterprises with purely private use cases may hire a neutral host vendor to deploy a cellular RAN for trusted internal use. This will provide business justification to invest in the small cell radios and related equipment, which can be extended to providing public cellular coverage in the future.

[4] https://www.phonearena.com/news/Cheat-sheet-which-4G-LTE-bands-do-AT-T-Verizon-T-Mobile-and-Sprint-use-in-the-USA_id77933

[5] One key technology for RAN sharing is called Multi-Operator Core Network (MOCN). https://www.gsma.com/futurenetworks/wiki/infrastructure-sharing-an-overview/

Convergence into a Single Physical Housing

A commonly question is whether multiple enterprise RANs will be converged into a single physical device. The answer is yes – and no. There are several competing dynamics that must be balanced for RAN convergence in a single platform including:

- **Propagation.** RANs which use similar frequency ranges – and therefore have similar propagation characteristics – are good candidates for convergence. Examples would be Wi-Fi, Bluetooth and Zigbee in the 2.4 GHz band. However, devices with greater propagation – such as a licensed cellular small cell operating below 2 GHz – will almost certainly have to be available in non-converged formats. The same applies in the reverse. RANs with very limited propagation – such as millimeter wave 5G or Wi-Gig that may not pass through more than one or two walls – will also likely have to be sold in standalone formats.
- **Power.** Radios consume significant energy to operate, and modern MIMO radios with 2, 4 or even 8 'radio chains' increase this requirement linearly. Therefore, the Power-over-Ethernet (PoE) budget available from an access switch is the single biggest limiting factor after propagation. As of this writing, an 802.3at switchport delivering 30 watts can generally only power two full-featured 4x4 MIMO class radios, along with one or two smaller IoT radios. 60-watt PoE is starting to be adopted in the enterprise, but it is challenging to cost-justify upgrading many switches when just a handful on each switch require 60 watts.
- **Lifecycle.** In general, network infrastructure upgrades are driven by turnover in the end device population. Wi-Fi 6 devices require Wi-Fi 6 access points to use the latest features. CBRS or 5G devices require compatible base stations. And Wi-Fi and cellular technologies have completely different upgrade cycles. Given that different devices are upgraded on different time scales – especially in the retail or manufacturing sectors where ten-year lifespans are common for many device types – careful consideration should be given to converging dissimilar RANs in the same housing even where propagation and power do not pose a challenge.
- **Control plane.** Cellular systems use a completely different internal architecture and protocols than Wi-Fi systems. IoT radios like Zigbee are again different from Wi-Fi in key respects. Any hypothetical "super" cell containing all these radios would almost certainly require separate management platforms, each with partial responsibility for the device.

As a result of these factors, it is unlikely that a multi-band Wi-Fi AP will be combined with a multi-band small cell for an enterprise buyer. At most, we expect to see small cells with one single-band Wi-Fi radio – primarily to minimize cable pulls for neutral host small cell retrofits. As a best practice, we instead recommend engineering each enterprise RAN based on its specific requirements for coverage, device lifecycle, backhaul speed and other attributes.

None of this is to say that 'combo' units will not exist, but we believe that convergence is most likely to occur for enterprise RANs with similar propagation, power and lifecycle requirements.

CITIZENS BROADBAND RADIO SERVICE (CBRS)

As mentioned previously, CBRS is a recent development pioneered in the US. The FCC – which has been credited with kick-starting the Wi-Fi market by liberalizing unlicensed spectrum – is enabling 'three-tier spectrum sharing' in 150 MHz of spectrum from 3550–3700 MHz.

Spectrum sharing in this new band will allow existing incumbent users and other 'grandfathered' licensed incumbents to be protected, while new installations are allowed where they will not interfere with these protected incumbents. These incumbents occupy the first 'tier' and take precedence over the other two tiers, as shown in Figure 3. Tier 1 users are only found in specific locations, typically use only a small portion of the band, and their operations may be temporary / occasional in nature. Coordination between and within the three tiers is managed by third party spectrum managers that operate Spectrum Access Systems (SAS) that are certified and overseen by the FCC.

In the second tier, private organizations may purchase limited licenses called PALs (Priority Access Licenses) at auction to operate radios in the lower 70 MHz of this band. Because the minimum geographic area of a PAL is quite large – an entire county – it is expected that these will be primarily

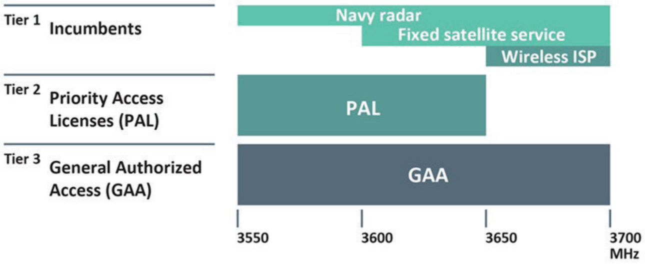

Figure 3: CBRS Band Tiers

purchased by existing MNOs. PALs are purchased on a county basis – and there are over 3,100 counties in the US. These licensees will have the exclusive right to use the RF channels they purchased after the installation of the base stations operating in the CBRS band, so long as they do not interfere with incumbents in the first tier. A third-tier of priority, GAA (General Authorized Access) allows any other private organization to opportunistically use CBRS channels wherever they will not interfere with incumbents or the PAL users, but they will have to vacate the channel if any higher-priority licensee starts transmitting nearby to particular CBRS access points.

In addition, there are certain constraints on access to spectrum that likely pose no operational challenge but should be fully understood before purchasing a CBRS system. The protection of Tier 1 incumbents – generally 5 – 10 MHz at a time per location – comprising military users (mostly mobile, long-range ship-borne radars), satellite base stations and others, employs regions called dynamic protection areas (DPAs) that on paper cover around 40% of the US population.

GAA deployments within 50 – 100 miles of major coastal cities including Boston, New York and Los Angeles are within range of these incumbents. Specialized equipment called an Environmental Sensing System (ESC) is used to identify and protect usage by tier 1 incumbents in each DPA. Each SAS operator has its own network of ESC sensors. Figure 4 shows the DPAs along the coastline of the continental United States. Each DPA has one or more ESC sensor that continuously listens for naval activity in the CBRS band.

If activity is detected, both PAL-based and GAA-based CBRS networks nearby may find their spectrum access temporarily switched to another channel by the SAS. Such changes must occur within 5 minutes of detection. Theoretically, under extraordinary conditions, GAA operations may be terminated temporarily if there is not enough spectrum to accommodate both the channels in use by the Navy and the purchased PALs in that area. As a practical matter, this risk is negligible for several years while CBRS deployments are ramping up and there is plenty of spectrum. Enterprises with a need for guaranteed access to CBRS spectrum are eligible to bid on PAL licenses.

PAL channels in a specific area may also be used by GAA access points in two cases. First, as noted earlier, because the three-tier system operates on a “use it or share it” basis all PAL spectrum is usable for GAA operations unless and until the PAL owner deploys equipment in that location. Second, PAL spectrum can be subleased from its owner. While each PAL covers an entire county, it is entirely possible that the licensee only intends to deploy equipment in a

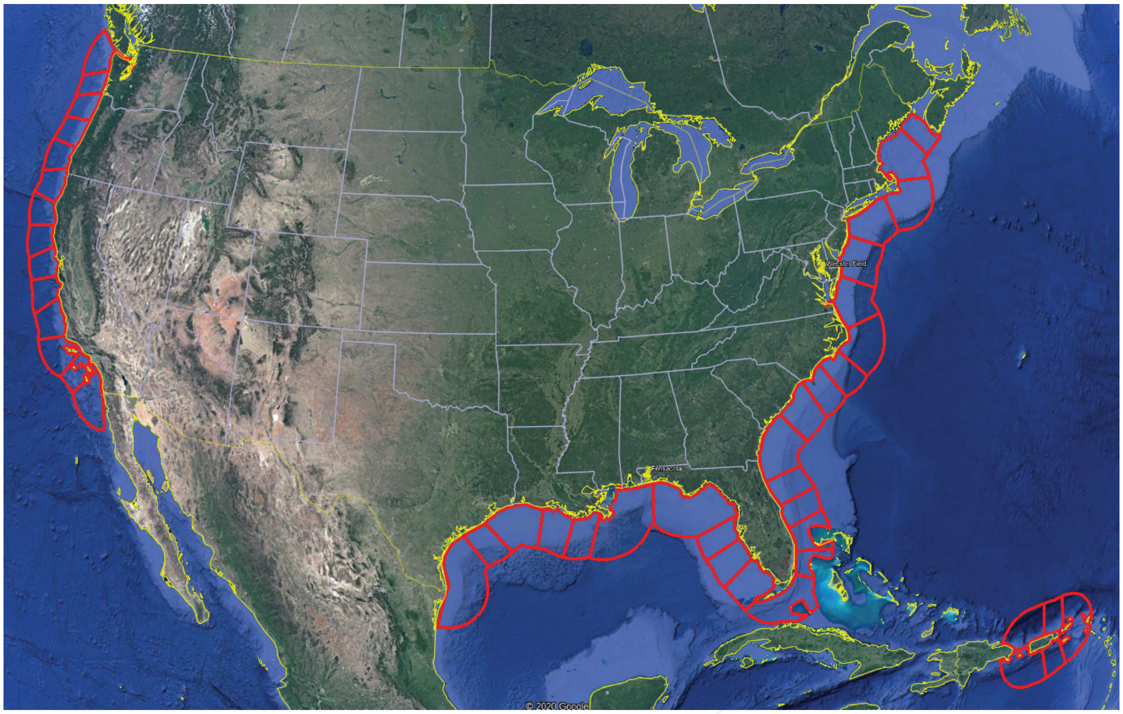

Figure 4: CBRS Dynamic Protection Areas (DPAs)

portion of the country. For example, an MNO that acquires a PAL may only be interested in deploying small cells in hot spots such as malls, stadiums or mass transit facilities. The rest of the county may therefore be available for subleasing. These transactions are permitted under FCC rules and are known as *secondary market licenses*. SAS operators are expected to facilitate such transactions. For an enterprise that desires the certainty that a PAL provides, but has no need to purchase a county-wide license, a secondary market PAL sublease may be an attractive option.

One key function of a SAS is to manage the boundaries between CBRS small cells. In fact, the SAS is responsible for assigning channel numbers to each PAL and GAA user. Each SAS maintains a heat map of all the CBRS devices in the US, and each SAS is contractually obligated to ensure a minimum level of protection between adjacent PAL holders. (GAA devices are not entitled to this protection.) Therefore, if a PAL holder does not intend to serve the entire geographic area to which it has purchased a license, it may notify the SAS that it may sublease those areas. This is called a *secondary market* and is a unique feature of CBRS that holds great promise.

The FCC authorized GAA operations in late 2019.[6] The PAL auction was successfully completed in August, 2020 raising nearly $4.6 billion.[7] Initial commercial deployments are under way by both private enterprises and MNOs (using Options 1 and 2, respectively). New devices supporting B48 are announced almost every month. Equipment manufacturers, neutral host providers and MNOs are collaborating in the CBRS Alliance to address the roaming challenges discussed earlier. The interoperability certification developed by the CBRS Alliance is called OnGo™.

WI-FI CALLING & ARUBA AIR PASS

Another way to improve the in-building coverage experience for public cellular subscribers is to leverage Passpoint-enabled Wi-Fi Calling using the Aruba Air Pass™ service. Aruba has partnered with major MNOs to create a cloud-based authentication service called Air Pass that enables mobile devices with SIM credentials to automatically detect and authenticate to participating customer WLANs with no action required by the user. Air Pass can be deployed by itself or in combination with a CBRS neutral host solution to reach even more subscribers.

Air Pass is a novel approach to the in-building cellular coverage problem that allows an enterprise to provide a dependable experience for any cellular mobile device. As noted above, *Wi-Fi in combination with Passpoint is functionally equivalent to small cell deployment Option 4 because multiple services and operators can be converged into a single radio layer*. In North America, Passpoint is supported by all nearly all major cellular operators. It can be piggybacked onto an existing Aruba Wi-Fi network via simple configuration changes and adding a secure RadSec authentication link from your facility to the Aruba Air Pass cloud service. Virtually every smartphone already supports this capability, although the user must manually enable Wi-Fi Calling before it can be used. Once activated, smartphones automatically connect to the WLAN and thereby gain access to data, Wi-Fi Calling, and SMS messaging services.

In this section, we will review how Passpoint works, how it compares to the cellular RAN techniques previously considered, and finally some of the challenges still to be solved. Wi-Fi, unlike DAS or small cells, is inherently multi-operator and Passpoint profiles are found on most SIM-based devices. Experience at Aruba trials and customer events has shown that enabling Passpoint results in an immediate connection for most of the devices that the public carry in their pockets with no action by the user.

Wi-Fi Calling Adoption

Wi-Fi Calling (WFC) allows MNO subscribers users with a compatible smartphone device to make and receive voice calls over Wi-Fi, even in "airplane" mode. WFC is generally also combined with SMS and text messaging service. WFC is now pervasive globally – with over 127 operators in 47 countries active as of June 2019.[8] In North America alone, WFC is supported by 30 operators. Europe has 49 operators that support WFC, while the Asia-Pacific region has at least 24 operators. Nearly 10 operators in Latin America currently support WFC.

AT&T has led the industry as single most assertive adopter of WFC and of Passpoint to get their subscribers onto Wi-Fi networks, with other operators now working to catch up. By leveraging the widespread availability of WFC,

[6] https://www.cbrsalliance.org/news/fcc-authorizes-historic-ongo-deployments-in-3-5-ghz-cbrs-band-opens-billions-in-economic-opportunity-for-u-s/

[7] https://www.fiercewireless.com/regulatory/cbrs-3-5-ghz-auction-concludes-raising-4-58b

[8] https://support.apple.com/en-us/HT204040

enterprises that use Passpoint to seamlessly steer devices onto a suitable Wi-Fi network can offer a comprehensive, multi-operator solution for on-campus and in-building cellular coverage.

What is Passpoint?

To use Wi-Fi to provide robust in-building or campus cellular coverage requires Passpoint, which enables a smartphone to automatically discover and attach to an authorized Wi-Fi network. When Passpoint is enabled for Air Pass, the WLAN advertises the PLMNIDs of participating MNOs. These advertisements are detected by the phone and trigger the phone to connect to the specified WLAN automatically.

Passpoint traces its origins to the IEEE 802.11u amendment (also known as 'Hotspot 2.0') and subsequent Wi-Fi Alliance (WFA) 'Passpoint' certification program. Following the initial Passpoint launch, the Wireless Broadband Alliance (WBA), with Aruba participation, conducted a series of international "Next Generation Hotspot" carrier trials to verify performance in real-world networks and applications. Thus, this technology is well-understood and established. Over the last few years, it has been adopted by AT&T Wireless as its preferred protocol for national Wi-Fi offload and international roaming, and it is supported by all major North American service providers. Internationally, operators in Europe and Asia are beginning to deploy Passpoint profiles as well.

Passpoint offers possibilities beyond local cellular offload. International roaming for many major operators can be supported, a growing community of smart-cities is already starting to use Passpoint to offer roaming across the world, and enterprises can configure their own Passpoint profiles on devices for private guest-access purposes.

How Do Passpoint and Air Pass Work?

A Passpoint-enabled cellular offload service requires three functions depicted in Figure 5. Smartphones must be Passpoint-capable and configured with profiles that identify service providers; Wi-Fi access points must advertise a list of supported service providers; and a secure link must be established to the roaming exchange for authentication.

For the first component, configured smartphones, the popular smartphone families (Android, Apple iOS and Microsoft Windows 10) have incorporated Passpoint capabilities for some years now. All that remains is for mobile operators to specify Passpoint-certified devices, and to add the appropriate configuration profiles so that each profile contains information identifying the service provider. For example, a smartphone could identify an AT&T-capable access point by:

- 'MCC-MNC' (Mobile Country Code – Mobile Network Code, e.g. 310-410)
- 'NAI realm' (e.g. attwifi.com or wlan.mnc410.mcc310.3gppnetwork.org)
- 'OI' (Organization Identifier, not widely used)

The profile for the MNO will contain at least one of these elements together with login credentials which for most cellular operators will point to the device's SIM card. The smartphone OS vendors have added functionality to allow network-based configuration of Passpoint profiles based on clicking links. But mobile operators prefer to configure Passpoint profiles on the special software builds they load on or push to subscriber devices.

Once a smartphone has a Passpoint profile, it regularly scans for Wi-Fi access points that use Passpoint to advertise their service provider reachability. This is done before association, so the smartphone can pick a suitable access point and know it will be able to get service before initiating authentication. The WLANs of Aruba customers that have enabled Air Pass are already advertising the network identifiers of our MNO partners.

One key advantage of the Passpoint profile is that it is not linked to a particular SSID, so a single client profile will work across any WLAN that has appropriate Passpoint configuration. Organizations can add Passpoint service to any existing SSID that uses WPA2 or WPA3 authentication.

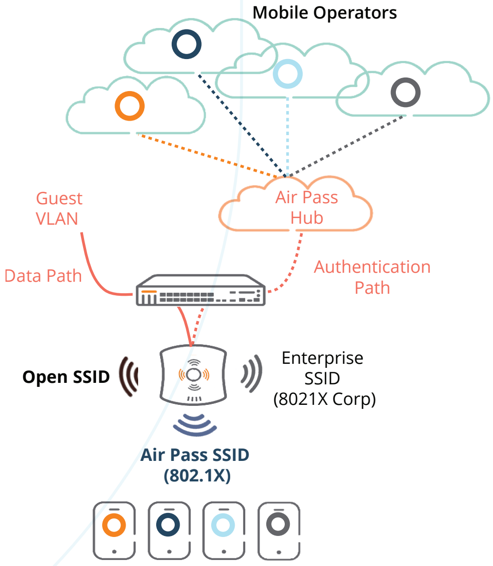

Figure 5: Air Pass & Passpoint Implementation

Once the device has detected a match between a pre-configured profile and an access point's advertisements, it starts to authenticate. For this, it needs a path to the MNO's home subscriber server, which is provided by a RadSec connection from the Aruba Mobility Controller to the Aruba Air Pass roaming hub.

Authentication protocols supported by Passpoint include:

- EAP-SIM/AKA/AKA', using SIM credentials
- EAP-TTLS, using username-password
- EAP-TLS, using X.509 certificates

All these options use the WPA2/WPA3-Enterprise 802.1X protocol which protects credentials and completely encrypts over-the-air traffic.

The cloud RadSec link is used only for authentication traffic following the 802.1X protocol. Data-plane traffic remains on-campus and is routed by the WLAN. Most Aruba customers place Passpoint-attached clients in a guest role on the existing guest VLAN. This ensures that offload traffic is completely segregated from the corporate network, leveraging Aruba's built in role-based access control capability. Aruba's experience suggests that for most traditional office environments, enabling Passpoint-based cellular offload adds a minimal amount of new traffic, even while greatly increasing the number of devices attached to the WLAN.

Air Pass Benefits

The combination of the Aruba Air Pass service, Passpoint authentication and Wi-Fi Calling enables robust in-building and campus cellular coverage, delivered over Wi-Fi. For *cost-conscious enterprises that do not wish or cannot afford to invest in a privately-owned cellular network like CBRS, this is a compelling solution*. It allows an IT department to extract more return on investment from an existing WLAN and can address coverage problems and increase capacity with minimal additional investment.

Because Air Pass dramatically increases the attach rate of smartphone devices to the WLAN, the utility of the Wi-Fi infrastructure as a sensor system is also enhanced. Applications such as shopper analytics in the retail vertical, space and energy/lighting optimization for facilities departments, and network security systems have greater visibility of visitor data traffic, location, and behavior.

Combining Air Pass with CBRS

Both Wi-Fi calling and an Option 4 small cell deployment using neutral host CBRS offer different paths to improving public cellular service on premise. However, as noted earlier, each of these solutions has certain limitations. It is possible to co-deploy them to maximize the number of mobile subscribers benefitted by these solutions.

We have seen that WFC requires a user to manually activate the Wi-Fi Calling feature on their smart device to supply emergency 911 information. Therefore, while Air Pass will connect all the smartphones in your building virtually to the Wi-Fi network, there is no guarantee that WFC will be available for everyone. All users still enjoy greatly improved data offload – which can improve cellular voice all by itself by reducing the load on the outdoor cell tower. Another constraint to be aware of is that not all users enable Wi-Fi.

Since CBRS-enabled phones began shipping in 2019, it will take a few years to reach a majority in most indoor environments. Also, CBRS is a US-only solution today whereas WFC is available worldwide. Finally, we noted that an Option 4 small cell system requires complex back-end roaming agreements and technical integrations with major MNOs. Adoption of MOCN is limited as of this writing, and the CBRS community has further work to do before it can be widely deployed.

The best solution for an enterprise with sufficient resources may be to deploy both Air Pass and either an Option 1 or 4 small cell solution. This offers enterprise network architects significant new tools to address business needs.

CELLULAR DATA PATH CONVERGENCE WITH ENTERPRISE NETWORKS

Enterprises with wide-area connectivity requirements have been implementing overlay networks for many years. These were depicted on the right side of Figure 1 and are typically accomplished with VPNs to impose a trust layer. With the advent of affordable privately-owned CBRS systems, the enterprise network architect and information security teams naturally wish to understand how cellular devices will integrate with the longstanding layer 2 / 3 segmentation strategies in use in their core and distribution networks. Today, the state of the art is relatively primitive by enterprise standards, essentially just raw forwarding of source-NATed cellular traffic. Identifying individual user equipment (UEs) to apply per-device policy inside the trusted enterprise perimeter is not possible with current technology from traditional cellular equipment makers. This in turn will force very coarse-gained network security strategies, such as segmenting off all cellular traffic into untrusted subnets for the foreseeable future. In this section, we will review the basic landscape of cellular data path integration as it presently exists.

Public Macro Network Topology for Enterprise

We begin with a look at the familiar wireless WAN "public" macro network in use today by some enterprises. The MNO issues subscriber identity module (SIM) cards to enterprises, who purchase voice and data subscriptions in bulk from the operator. The enterprise inserts SIM cards into compatible devices (called UEs in the cellular word, short for user equipment). These devices may roam anywhere that operator has coverage. This includes traditional mobile use on the "outdoor macro" network. It also includes a plethora of "indoor macro" technologies ranging from small cells to distributed antenna systems (DAS).

Figure 6 takes a technology independent view of this, without regard for the specifics of 5G or LTE system architectures. The only assumption below is that the SIM is issued by an MNO. This topology represents a distributed WWAN operated by a single MNO in a geographic region – a single state/province, a single country, or even a multi-country network such as exists in Europe. By definition, this use case involves a business requirement for large area coverage, such as connecting fixed vending machines across a city or connecting private automobiles in motion across a country.

In this simplified topology, traffic from enterprise devices being served over the WWAN traverses the MNO core network, where it is routed to the public Internet via the nearest access point name (APN), and then on to enterprise compute resources in public or private clouds that are reachable via public IP addresses. The APN can be thought of as the "border" between the trusted operator core and the external untrusted network. Each user device is allocated its IP address by a network element inside the operator core. From a user data path perspective, the public cellular topology is therefore an overlay network. Traditionally, a private datacenter deployment requires a VPN client on the UE.

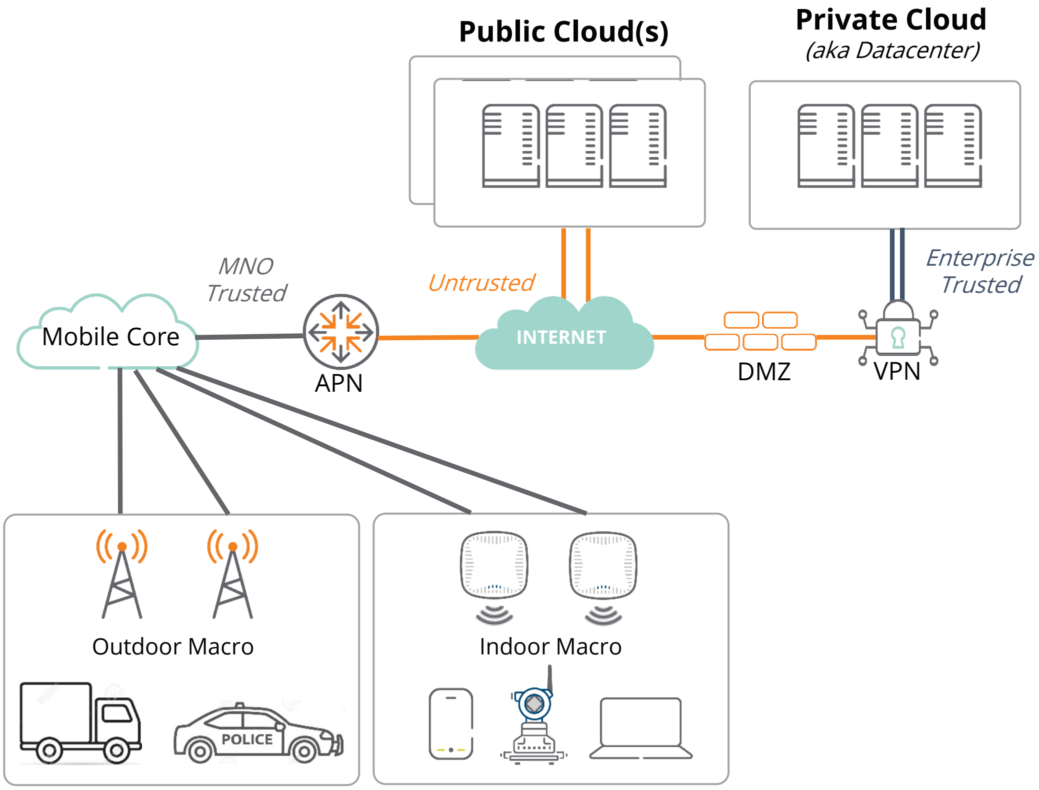

Figure 6: "Public" Cellular Network Topology for Mobile Broadband & IoT Devices

Operator core networks have historically been closed topologies or "walled gardens," especially 4G core networks. The original 4G architecture is well over 10 years old and was intended to be deployed by MNOs for their own use. There was no need to support multiple "tenants" on the network in the sense that it is understood by enterprise network architects, where different logical segments serve different user communities. In a 4G core network, the tenants are all subscribers of the same service. While there is rich support for billing flexibility to offer different pricing plans, in general all the traffic from all the users is co-mingled at L2/L3 because nobody else but the MNO has access to the system.

The 5G system architecture was explicitly designed with multi-tenancy in mind via a technique called network slicing. A network slice is analogous to a private MPLS service – it is a logical network segment with its own speed, latency, redundancy and other characteristics including the ability to route traffic. However, it is vital to understand that user plane traffic may or may not be isolated from other subscribers on a 5G slice.

Private Cellular Network Topology for Enterprise

As noted earlier, any privately-owned cellular network – whether 4G, 5G or a future 'G' is distinguished first and foremost from a public network by who issues the SIM credential. By definition, a private network uses privately issued SIMs. These SIMs may be physically provisioned and managed by the enterprise itself or by a service provider on behalf of an enterprise, but they belong exclusively to one entity and form a closed pool of secure credentials.

This single fact profoundly alters the network architecture as compared with a public macro network. Figure 7 illustrates a simplified view. As above, we initially consider this in a technology neutral manner without regard for specifics of whether the system is 4G or 5G.

3GPP based systems require a "mobile core" to function. The simplest way to think about this is that it is the operating system of a cellular network. This should not be confused with an enterprise network core which operates very differently. In the public model, we could safely ignore all the components of the mobile core because all its complexity including the provisioning and management of SIMs and radios is hidden from view and provided by an MNO.
In the private RAN scenario, the enterprise now must provide its own mobile core. There are two possibilities:

1. **Cloud core from service provider**: Just as the name implies, this is a third party hosted mobile core that is operated in a SaaS model. The term is misleading, since the radios are not in the cloud. Cloud RAN providers also typically provide SIM cards for a fee.

2. **On-premise core managed by enterprise**: This is the do-it-yourself solution. The enterprise purchases a mobile core software license from a compatible vendor and stands up its own server in a data center. It is analogous to putting up a unified communications server.

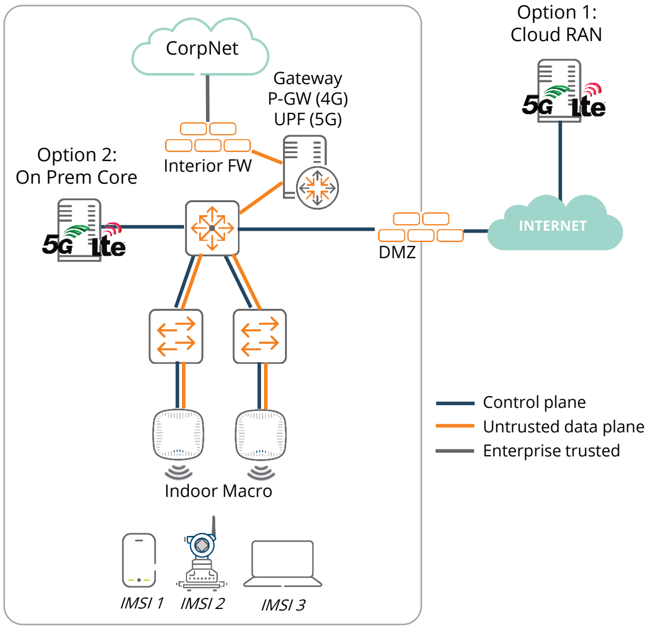

Figure 7: "Private" Cellular Network Topology for Mobile Broadband & IoT Devices

Cloud core is expected to be the dominant consumption mode for private RANs because the 3GPP architecture, operational concepts, and especially configuration model is extremely complex and has no relationship whatsoever to enterprise networking. Both options are shown in Figure 7, and the blue lines indicate 3GPP control plane traffic between on premise radios and the core.

The user data path – shown by the orange lines – is more interesting. In the "on premise" breakout topology, traffic is tunneled from the radios to a gateway which runs in software on an X86 platform. This is very similar to how controller-based wireless LANs (WLANs) tunnel traffic from a Wi-Fi access point (AP) to a Wi-Fi controller in a data center. The gateway effectively source-NATs this traffic to an untrusted L3 point-to-point link to an interior firewall on the corporate network.

Implications for Enterprise Network Integration

Under the 3GPP architecture, it is not possible to bridge traffic directly from a cellular device to an Ethernet network. This is because 3GPP devices do not have compatible MAC addresses. As you may be aware, cellular devices are identified by the combination of a unique hardware identifier (the International Mobile Equipment Identifier or IMEI) and a unique subscriber identifier (the International Mobile Subscriber Identity or IMSI). The gateway thus performs L2 address conversion, stripping the 3GPP headers from the L3 payload and inserting an L2 Ethernet address. It must also NAT the L3 traffic since the IP address space used inside the mobile core does not exist in the enterprise network. The collective term for this function in a 3GPP system is 'local breakout'.

It should be apparent to an enterprise network architect from the foregoing description that private cellular traffic must be treated in bulk on a segment basis, and that per-UE forwarding policy is not possible. This is due to three factors:

1. The source-NAT conceals the originating IP address and port number of the UE making it impossible to differentiate L3 sessions from different devices
2. Cellular UEs have no IEEE-compatible L2 identity
3. Enterprise firewalls and AAA servers cannot currently communicate with mobile cores to establish identity binding tables

As a result, L2 access control lists are the best that can be expected. Per-user role-based access control is not possible with these systems. Even this is problematic for some common enterprise network services. For instance, common enterprise peripherals that use L2 resource discovery protocols like mDNS are incompatible with 3GPP devices, which do not understand IEEE L2 frames. So, Apple TVs, Chromecasts, printers and some file sharing services will be unavailable to all cellular UEs.

All Cellular Traffic Must Be Segmented and Untrusted

Imagine if you had a neutral host provider that was supporting the ability for public SIMs to roam onto your privately-owned LTE network with access to your corporate network. How does the corporate network know what security policy to apply to traffic from that device? It cannot because the authentication systems are not compatible. There is no way to associate a particular IMSI+IMEI combination with an enterprise user in the AAA server. The only solution is to treat all such traffic as untrusted and place on the "guest" VLAN.

But this problem is not limited to public networks. It also affects purely "private LTE/5G" networks for the same reason. As of today, privately issued SIMs also cannot be "bound" to enterprise AAA identities. Cellular devices lack MAC addresses and, in any case, will all share the common MAC address of the local breakout device.

This means that local breakout for private cellular RANs must terminate in either (1) an isolated VLAN that is explicitly trusted with direct access to any required network services and firewalled from any other corporate network segments; or (2) an untrusted VLAN with at most access to the Internet but no other corporate resources.

The closest analogy to this with which many enterprises are familiar are Payment Card Industry (PCI) requirements for WLANs that process credit card payments. Individual credit card scanners often use preshared keys and cannot be individually identified over Wi-Fi, and Ethernet-based scanners have no authentication at all. These devices have been frequent successful targets of hackers. As a result, they have been required to be completely segmented onto isolated and deeply untrusted L2 VLANs – even though the devices are owned by the enterprise and nominally under enterprise control.

Until a private LTE product is available that bridges these fundamental architectural limitations in the 3GPP system design, role-based access of both public and private SIMs is not feasible in enterprise networks. However, using established segmentation strategies for untrusted devices, they can nevertheless be integrated with the enterprise data path in a controlled manner.

5G IN THE ENTERPRISE

The reason MNOs and their technology suppliers are so excited about 5G is because they view it as an enabler to enter new markets and create new revenue streams. These include consumer connectivity (using fixed wireless access or FWA) and device-centric use cases, such as industrial IoT applications. They also will use 5G to target enterprise connectivity, with some going so far as to call 5G a "Wi-Fi killer" and the "next generation local area network."[9] These new revenue streams would be timely, as carriers are trying to offset strong business headwinds, such as declining revenue per user in their core cellular businesses, a dwindling landline user base, and increasing competition from non-traditional service providers, such as Amazon and Google. These trends have resulted in rising corporate debt loads and poor stock performance. Construction of the 5G networks themselves will be extremely expensive, with estimates of as much as $275 billion dollars plus tens of billions of additional spending to purchase 5G spectrum in the 3 GHz and millimeter wave bands.[10] The carriers need to act.

Operator 5G Will Not Become a Next-Generation WLAN

We looked closely at 5G and compared it to Wi-Fi and other wireless access technologies to determine how each best serves our customers. 5G and Wi-Fi 6 represent different approaches to wireless connectivity. Yet both are based on the same technological building blocks (e.g., OFDM, MIMO, and higher-order modulation). 5G and Wi-Fi 6 improve upon the performance and economics of LTE and 802.11.ac, respectively. Like all cellular technologies, 5G fits when the user requires macro coverage and mobility and can afford to pay the additional cost for these capabilities. Wi-Fi 6 has excellent in-building mobility but does not roam well at high speeds. Importantly, Wi-Fi 6 is at least on par with 5G in terms of throughput, latency, spectral efficiency, reliability and connection density. From a security perspective, 5G finally catches the cellular network up to EAP based enterprise-grade encryption and mutual authentication systems that have been in use for years.

One of the most important differences between 5G and Wi-Fi is economics. To succeed as a next-generation WLAN in the enterprise, carriers would need to offer 5G service at a competitive cost to Wi-Fi. Earlier, we noted that 4G macro cellular technology does not adequately penetrate buildings to provide enterprise-grade service and that the "core" 5G spectrum in the mid-3 GHz range will magnify this problem. Therefore, to serve the enterprise market with 5G, some MNOs are proposing that customers deploy (and pay for) DAS or small cells to extend macro coverage indoors (in an Option 2 or 3 mode). DAS and multi-layer small cell systems are substantially more expensive than Wi-Fi systems from both per-square-foot and lifecycle perspectives.

The cost to implement 5G-based access in the enterprise does not stop with the cost of the network itself. Cellular service requires all laptops, printers, Apple TVs, and other connected devices to contain 5G-compatible cellular modems. These modems cost tens of dollars per device wholesale, and typically more than $100 to an end user. Every device also needs to be included in a service contract. Since almost none of these devices contain cellular modems, they would need to be upgraded or outfitted with external dongles. We doubt many enterprise customers will be willing to pay this additional cost and replace most of their equipment for an unclear benefit. Research analysts that track radio chipset shipments agree.[11]

Public 5G and Wi-Fi Are Better Together

The Wi-Fi *versus* 5G narrative misses the point. Wi-Fi and cellular technologies are both continuously evolving to better serve end users, and both markets will grow to serve the macro trend of connecting devices. Wi-Fi will continue to prove its value as a reliable, secure and cost-effective wireless access technology for most enterprise applications as it does today. In fact, the most likely outcome is that MNOs will come to depend on Wi-Fi more than ever, as operators begin to confront the through-wall propagation challenge with mid-band and millimeter wave 5G spectrum. The obvious solution to this problem is to leverage a neutral host technology that is already ubiquitous on the ceiling indoors and that is funded and upgraded regularly by the enterprise – namely Wi-Fi.

Indeed, the 5G standards contain vital innovations that make it possible to opportunistically bond an outdoor macro 5G network to an indoor Wi-Fi network to provide a high-quality device experience. For 5G and beyond, 3GPP decided to decouple the evolution of RAN technology from that of core network technology. A by-product of this decision is that the

[9] https://www.telecompetitor.com/donovan-outlines-att-5g-enterprise-strategy-company-plans-nationwide-5g-for-first-half-of-2020/

[10] https://www.accenture.com/us-en/insights/strategy/5g-network-build

[11] https://www.abiresearch.com/press/wi-fi-retain-connectivity-crown-5g-era-wi-fi-6-chipset-shipments-break-1-billion-unit-barrier-2022/

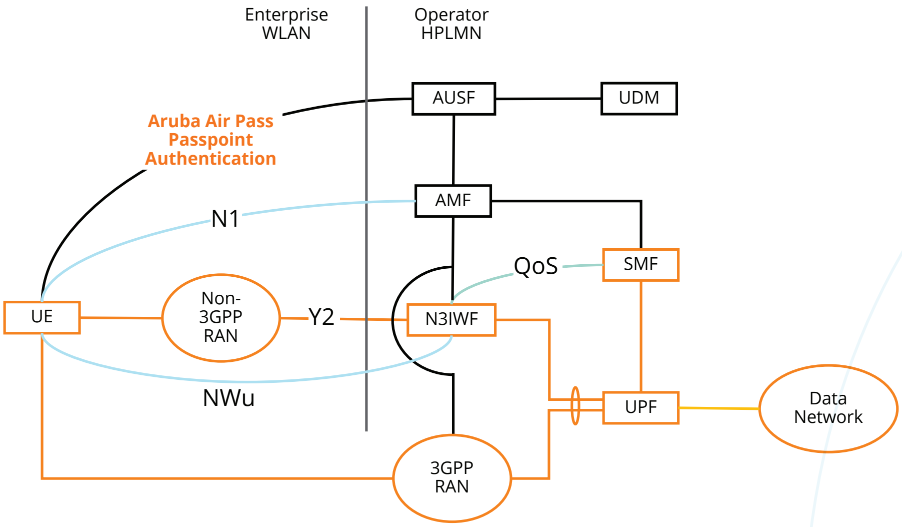

Figure 8: Aruba Air Pass Provides a Bridge Between 5G and Wi-Fi in the Enterprise

5G core is now RAN-agnostic and will manage both 3GPP and non-3GPP radios under a common architecture. This framework extends all the way down to such areas as having a common cryptographic key hierarchy that spans both Wi-Fi and 5G radios.

Aruba is investing in capabilities that make Aruba Wi-Fi the preferred partner for 5G networks. This includes a new feature in AOS 8.6 called Air Slice which replicates the guaranteed QoS capabilities promised (but not yet delivered) by 5G. Perhaps the most strategic initiative is Air Pass, which is designed to harness the RAN agnostic 5G core to enable transparent bonding of participating Aruba customer WLANs with participating 5G macro networks. Figure 8 depicts a simplified 5G core network architecture, showing the vital role that Air Pass plays in enabling dual mode 5G / Wi-Fi mobile devices to discover and opportunistically aggregate capacity from both systems.

This architecture – once fully realized – will take the concept of the multi-RAN enterprise to an entirely new level by incorporating third-party RANs into the broader access layer strategy of the enterprise network. This is what we mean when we say that Aruba believes that there is no one-size-fits-all answer to access-layer connectivity at the edge. Aruba is truly committed to a holistic approach that integrates cellular and non-cellular technologies over time. For Aruba customers that do not wish to or cannot justify the cost of privately-owned cellular infrastructure, this strategy provides a concrete roadmap to getting the best of both worlds. For MNOs, Air Pass offers a compelling way to work cooperatively with the enterprise as a roaming partner gain access to over 20 billion square feet of high-quality indoor WLAN coverage without spending precious capital equipment funds.

Public cellular connectivity will become another mission-critical service on top of enterprise networks, much as IP telephony did nearly two decades ago. Faced with 5G subscriber complaints and continued enterprise resistance to investing in small cells, mobile operators will partner with enterprise network vendors like Aruba to facilitate seamless hand-in / hand-out while leveraging 5G's unique capability to combine dissimilar RANs into a single capacity pool. And every dollar that enterprises invest in their Wi-Fi infrastructure is also a dollar invested in 5G readiness.

Private 5G and Wi-Fi Are Better Together

Enterprises that do have a business case for deploying privately-owned LTE networks such as CBRS are in a very enviable position. Once affordable enterprise-grade 5G small cells become available, the same RAN bonding capability can be deployed inside their facilities. This will allow dual-mode devices to simultaneously push traffic over both shared CBRS and unlicensed spectrum. The 5G core network, the WLAN controller and the mobile device will work together cooperatively to maximize the performance of each layer 3 session.

Neutral host integration of privately-owned 5G networks in the Option 4 scenario also gets simplified. As we documented earlier, the 5G core network is extremely rich in functionality as compared with its predecessor. This extends to integrating third party owned networks.

That said, cellular equipment manufacturers are concentrating on meeting demand from their MNO customers who are building out their 5G macro networks. Affordable enterprise-grade small cells are not expected in volume until 2022. In the meantime, enterprises considering an investment in LTE over CBRS should be certain to plan for a migration strategy. It may even be the case that there will be no need to upgrade an enterprise LTE system for several years given that 5G client devices will likely take some time to materialize in all the form factors of interest for dominant use cases.

CONCLUSION

The multi-RAN enterprise has been a reality for several years, with Wi-Fi and Bluetooth and Zigbee prevalent in all manner of organizations worldwide. Seen against that background, the emergence of affordable privately-owned cellular RANs is simply the latest tool available to enterprise network architects to address particular business use cases. In the United States, CBRS shows great promise at simplifying what until now have been the twin intractable problems for private LTE: access to spectrum and consolidating multiple MNOs into a single layer infrastructure. Depending on the use case(s) to be served, cellular RANs can be deployed independently of Wi-Fi, or may be co-deployed with coverage engineered to match the Wi-Fi footprint to support client devices with both types of radios. In no case will either 4G or 5G systems replace Wi-Fi in the enterprise.

The aggressive adoption of Passpoint by North American MNOs has led Aruba to introduce the Air Pass roaming service, which enables the bulk of enterprises that have no need for private cellular systems to enjoy the benefits of automatically attaching subscribers of participating MNOs. This vastly reduces friction for BYOD devices and unlocks important new enterprise use cases for location analytics, smart spaces, and location-based network policy.

Over the next few years, the most promising enterprise uses of cellular RAN technology are for private voice and data applications. This includes mobile point-of-sale, IoT, push-to-talk voice and warehouse automation. In these scenarios, the enterprise owns the entire end-to-end system. Providing public voice roaming services between macro cellular networks and privately-owned enterprise RANs on a neutral host basis is challenging as of this writing. There are business constraints – primarily the establishment of roaming agreements – as well as technical challenges ranging from handover signaling to establishing the necessary secure connections to operator core networks. Solutions are being worked on collaboratively within the CBRS Alliance and mobile operators.

Integrating cellular device traffic securely into the complex enterprise L2 / L3 network architecture is feasible today but has important limitations that must be fully appreciated by enterprise architects and information security staff. The only practical strategy with existing platforms is to fully segment traffic from both private SIM and operator/public SIM devices onto untrusted L2 segments using the same DMZ techniques commonly applied to IoT devices and credit card processing equipment.

The transition to 5G will bring new flexibility for both enterprises that rely exclusively on Wi-Fi and those that have deployed privately-owned cellular over some or all of their facility footprint. 5G and enterprise network technologies can be integrated in a variety of ways – from loose coupling using an Air Pass solution to a tighter integration where the enterprise appears as a normal visited network. 5G offers the promise of creating a new and mutually beneficial relationship between mobile operators and enterprises, and effectively peering their networks under dynamic policy control.

aruba
a Hewlett Packard
Enterprise company

WP_CBRS5GandWi-Fi_SK_102020 a00105590enw

Contact Us **Share**