WHITE PAPER

SOLVING THE INDOOR WIRELESS COVERAGE PROBLEM: ALTERNATE APPROACHES

SOLUTIONS FOR ENTERPRISE

a Hewlett Packard Enterprise company

TABLE OF CONTENTS

This is a companion document to the *Solving the Indoor Cellular Coverage Problem* white paper.

INTRODUCTION

Despite many advances in enterprise and campus communications in recent years and the extensive build-out of cellular networks, enterprise customers consistently suffer from poor on-campus and in-building cellular coverage. They are seeking a simple, inexpensive, universal solution but none is currently available. However, recent developments in small cells, emerging 5G cellular technology and the new CBRS (Citizens Broadband Radio Service) band all offer hope that a more attractive technology will emerge in the near future. Meanwhile, cellular operators and the Wi-Fi industry have made significant advances with a solution called NGH (Next Generation Hotspot). In this paper we survey established and emerging solutions for in-building and on-campus cellular coverage, and contrast them with the Wi-Fi-based NGH, which is powered by Wi-Fi Calling and a technology called Passpoint.

The reasons for poor on-campus cellular coverage are well-documented. While operators are able to claim that over 99% of the population is covered by their networks, a surprising number of businesses, universities, even hospitals find themselves in that missing 1% and have little-to-no coverage from at least one of the major operators.

This highlights that cellular coverage is an overlay of multiple operators. If there are four major operators in a country and three have good coverage in a particular area, there is still a problem for subscribers of the fourth operator. Multi-operator installations are a requirement for any acceptable solution, but they are difficult to achieve at reasonable cost.

Meanwhile, in-building coverage from the outdoor macro network has deteriorated in recent years as newer technologies have moved up the radio spectrum: the higher 4G frequencies do not penetrate as far indoors as 2G and 3G. Even building construction is working against us, as the low-emissivity ('low-e') coated glass used to sheath most modern buildings becomes less transparent to cellular signals. The problem is becoming more acute.

As you may have seen deployed in large public venues, 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. There is a long-established movement looking for lower-cost alternatives to DAS.

Elsewhere, recent trends in the 5G upgrade cycle should help with on-campus and in-building cellular coverage. In order to support high data-rates at high frequencies, 5G networks will rely heavily on 'small cells': inexpensive, short-range base stations that can be easily and quickly mounted on poles, roofs and inside buildings. But small cells have been available for many years without achieving significant enterprise penetration, and their well-documented obstacles to deployment have not yet been completely solved.

Meanwhile the FCC (Federal Communications Commission) in the USA is in the process of opening up a new band at 3.5 GHz, for CBRS, based on an experimental 'three-tier licensing structure'. CBRS is intended to allow enterprises and organizations to set up private base stations for cellphones and other client devices inside buildings, however development has been sluggish. The FCC has (for the fourth time) revised the rules to make them more attractive to large operators, while device availability for CBRS is still unclear. Many non-US devices will not be supported, and the available spectrum for enterprise use is questionable.

Aruba has been monitoring small cell and CBRS technology for many years, but we have yet to see a widely-applicable, inexpensive enterprise solution incorporating both Wi-Fi (unlicensed) and cellular (licensed) radios in a single unit.

At the same time, we have observed that the cellular industry has quickly embraced WFC (Wi-Fi Calling) as an alternative coverage strategy, with 102 operators in 45 countries now offering WFC, including 9 in the US alone. WFC call quality is generally high, especially over managed enterprise WLANs which have been designed for voice-grade coverage for some years now.

Aruba has reached the conclusion that it is better to embrace Wi-Fi offload and roaming as a coverage solution for all devices like cellphones that incorporate both Wi-Fi and cellular interfaces. The seamless roaming mechanism of NGH uses a Wi-Fi Alliance certification known as Passpoint and is applicable to a large number of cellular operators and nearly all currently-available cellphones. In conjunction with the widespread availability of WFC, this offers a comprehensive, multi-operator solution for on-campus and in-building cellular coverage.

The remainder of this paper compares and contrasts these various approaches – of which none are perfect today – for private organizations and enterprises to offer on-campus and in-building cellular coverage.

REQUIREMENTS

Most of the requirements for on-campus cellular coverage are straightforward.

- **Inexpensive solutions:** Whereas in the past, cellular operators would take on some or all of the cost of providing indoor coverage, the recent trend has been to push these costs onto the end user organization. And the existing baseline solution, DAS (Distributed Antenna Systems) is expensive to purchase, install and operate.
- **Support for multiple operators:** An on-campus cellular solution must support all major operators in the geography. All the solutions listed here are capable of multi-operator configuration, but in most cases the cost escalates as more operators are supported, often in direct proportion to the number of operators. Each operator has unique technical performance criteria and operational procedures, so it is challenging both technically and operationally to build a single installation that satisfies several operators' requirements.
- **Support for as many devices as possible:** This includes as much of the installed base of cellular devices as possible. Along with the multi-operator requirement, this should ensure that nearly all of the devices consumers bring on-campus will be supported.
- **Support for the latest radio technologies:** New technology is clearly desirable, as informed consumers are aware that, emerging 5G technology is superior to 4G LTE, and 4G is superior to 3G. We can see the differences in application performance but DAS and small cells, with a limited number of internal radios, cannot span the cellular technology generations as well as the macro networks can. Sometimes, going with the latest technology means stranding older devices that may not support it (see support for as many devices as possible, above). For example, a single-radio small-cell can only be one flavor of 3G, 4G or 5G, and will be unable to support older devices. Also, leaky-coax DAS systems are inherently incapable of supporting MIMO (Multiple Input, Multiple Output, or multi-antenna) technology, an important performance-driver for modern cellular systems.
- **Clear demarcation of control and management responsibilities:** Since cellular frequencies are licensed to operators, DAS and small cell installations inside enterprise buildings and riding on their LANs, are usually managed and maintained by an operator, or a partner of the operator, to ensure that the transmissions in licensed bands are legal in terms of RF channels and power levels. Demarcation extends to who is authorized to install and maintain the small cells, where re-configuration and helpdesk calls are routed and how backhaul traffic is managed.

This is a short and non-controversial list, but no currently available solution meets all these requirements. Next we will examine the various options in more detail.

DISTRIBUTED ANTENNA SYSTEMS

The long-established DAS architecture is still the only solution, commercially-available today, that supports on-campus and in-building cellular coverage. DAS delivers the signal to in-building areas by placing antennas, sometimes 'leaky-coax' or small radios similar to Wi-Fi access points, around the building to communicate directly with cellular devices over short distances.

The antenna and radio units are connected, usually via proprietary cabling, back to a 'base-station hotel' in the basement, where one or more base-stations installed and managed by the cellular operators are mounted. Backhaul to the cellular network may be via an Internet connection but is usually over a dedicated communications link – operators are wary of sharing standard Internet connections with enterprise traffic.

Architecturally, each operator provides its own base station equipment, with the analog signal from its antenna connectors distributed inside the building by the DAS system - ensuring good coverage. Systems often include amplifiers and active electronics to condition and split the signal to multiple antenna units over coax or optical cables, which are installed above the ceiling and dedicated to the DAS, but from the operator's perspective the entire DAS is a passive element that relays their RF signal. Once a DAS system is installed, it continues to consume significant floor space and energy.

DAS Limitations

DAS is well-understood, and generally works well. But equipment is expensive and uses proprietary protocols, extensive in-ceiling installation work for cabling and interconnections is often difficult and expensive, and coordination with operators is cumbersome. Indeed, operators often refuse to allow their signals to be shared over distribution networks or amplifiers with other operators' signals, so a DAS for 4 operators may cost 4x a single-operator installation. For this reason, independent 'neutral-host' DAS operators have emerged in recent years, managing all this complexity for the enterprise customer but often adding to the cost.

It is also difficult to upgrade DAS to the latest cellular technology, so over time an installation will become outdated. Many older DAS installations still use 3G technology because of the high cost of upgrades to current standards, and as mentioned earlier, MIMO is only supported on by a very few DAS architectures. But a DAS can support nearly all cellular devices, and in addition, some non-cellular wireless networks.

DAS solutions are also the only viable option for venues that require very high-density cellular capacity, such as airports and stadiums. This is because a DAS can potentially carry every licensed frequency from every operator simultaneously over a common antenna system, enabling each operator to deliver maximum spectrum availability and network capacity.

Its most significant drawback is cost, especially where multi-operator support is required. For this reason, only a small percentage of buildings have a DAS system today.

SMALL CELLS

At its simplest, a small cell is a miniaturized base-station, capable of light-pole or indoor wall-mounting. It has a low-power radio, with a range of perhaps 100 meters, and lower capacity in terms of user-count, but in other ways the radio functions like a macro base-station. Depending on the application, its backhaul connection may be on fiber or over Ethernet.

The concept was first proposed around 2007, when the Femto Forum was formed. As markets developed and requirements changed, the Femto Forum changed its name to the Small Cell Forum in 2012. Enterprise-class small cells (which may also include Wi-Fi radios) will benefit from lessons learned in the build-out of macro 4G cellular networks, where they are being deployed by major operators.

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 vast number of small cells deployed inside numerous private facilities, but it introduces some challenges.

Perhaps the greatest challenge is multi-operator support. A single-radio small cell can only run on one frequency at a time, which limits it to a single RF channel, technology and operator unless roaming agreements, a neutral-host or infrastructure-sharing arrangement are in place. The Small Cell Forum lists several architectures that facilitate multi-operator sharing of a single-radio small cell, but so far, very few operators world-wide (and none in the US), have developed the technical and commercial agreements to allow such sharing. It remains to be seen whether this will change in the US market, but many analysts have remarked that this is an essential development before the in-building and on-campus small cell market takes off.

A second challenge stems from the arms-length arrangement most operators use for managing small cells. While small cells are deployed to counter poor cellular coverage, they are seldom out of range of the macro network. This means their operating frequencies need to be coordinated with neighboring macro coverage so they must take locally-unused frequencies, or reduce their transmit power. Interference between the macro network and small cells can be a significant problem, and solutions are still incomplete. It also means that handovers, both within the small cell network, and between small cells and the macro network must be carefully managed, a difficult problem that is still not comprehensively solved in today's equipment.

Small Cell Limitations

Since small cells have capacity limits in user-count and throughput, their usage is managed. Private organizations deploying small cells would prefer to reserve them for employees, but this is difficult for an operator to manage: how can it determine if a phone connecting to the small cell is an employee, a guest or a random consumer passing by the window? This question of which users are authorized to use a small cell still presents a challenge in some applications.

Ownership, management and control are also 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 walls of their buildings, back-hauled over their LAN but managed by a remote entity. This becomes a greater issue when re-configuring or troubleshooting.

Small cells are making steady progress in the macro network. As measured by unit shipments, the bulk of their market to date has been for macro network densification in urban and suburban areas. They can be deployed at low elevations below buildings and other ground 'clutter', and hence reuse frequencies while providing faster speeds. By virtue of stronger signals, they are closer to users. Small cells that operate in the 5 GHz band used by Wi-Fi are being deployed outdoors to enable 'gigabit LTE', using channel-aggregation techniques to increase network capacity.

Despite low market penetration and the objections above, small cells hold some promise for on-campus and in-building deployments. It is possible to describe a very comprehensive indoor small cell solution on paper, as the standards and architecture are well-defined, but time will tell whether these significant challenges can be overcome.

CITIZENS BROADBAND RADIO SERVICE

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 is the ability for different organizations to use the same frequency bandwidth in different location, and will allow existing government users and other 'grandfathered' licensed incumbents to be protected, while new installations will be allowed where they do not interfere with incumbents. This will allow private organizations, for example hospitals, enterprises and universities, to purchase limited licenses called PALs (Priority Access Licenses) to operate radios in this band, where the licenses will apply to an area of a few square miles and last for several years. Under the current proposal, these licensees will have the exclusive right to use the RF channels they are allocated by a SAS (Spectrum Access System), but only where and when they will not interfere with incumbents such as military users (mostly mobile, long-range ship-borne radars), satellite base stations and others.

A third-tier of priority, GAA (General Authorized Access) allows private organizations 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.

CBRS is an important initiative, and the FCC is to be congratulated for taking innovative steps to allow greater use of spectrum – a scarce resource – in the US. But the scheme is not without risks.

CBRS Deployment Risks

The first risk is regulatory uncertainty. The FCC has revised the rules for licensing the band four times over the last six years. The latest revision made significant concessions to operators' demands, increasing the coverage area per license, and extending the license terms in ways that make them less attractive to private organizations. It is clear that operators will continue to lobby to bring the spectrum into harmony with global standards, which designate it for fully-licensed cellular use. Meanwhile, the band was opened for GAA use in late 2019, with the first auctions of PAL licenses tentatively set for mid- or late-2020.

Secondly, the reliability of access to spectrum for GAA users is in question. The protection of incumbents uses exclusion zones that cover around 40% of the US population. In particular, GAA deployments within 50 – 100 miles of major coastal cities including Boston, New York and Los Angeles may find their spectrum access temporarily terminated by the SAS due to routine naval operations offshore. Spectrum that may be interrupted without warning is of limited use to organizations building communications systems.

A third business risk that must be considered is lock-in to a third-party neutral-host provider. CBRS will employ TDD-LTE (Time-Division-Duplex, Long-Term Evolution) radios to provide neutral-host cellular service. This means that a 4G core network (Evolved Packet Core, or EPC) is a necessary component of any CBRS deployment. However, few enterprises have the technical knowledge to deploy and operate an EPC. Therefore, the expected business model is for third-party service providers to deploy small-scale core networks on-premise, or offer 'cloud radio access network' (cloud RAN) services.. These providers will have roaming agreements with various cellular operators, and the enterprise network will appear like a separate roaming network to consumer devices. But it also means that changing neutral-host providers in the future will be complex and expensive.

CBRS End-user Device Risks

At the device level, other challenges are clear. The first is device availability: how soon will the devices enterprises care about – cellphones, tablets, PCs, and others – be modified to support the CBRS band? High-end smartphones are beginning to add support for LTE channel 48 (the CBRS band), but it will take many years for the installed base of consumer devices to turn-over so that most consumers, if that is the desired target audience, carry a CBRS-capable device. Another challenge is that CBRS is a US-only initiative, so organizations with many non-US visitors may not find it a comprehensive solution.

CBRS, while a laudable concept, has not proven itself capable of moving from slide decks to enterprise networks. If it does, then it will encounter the limited range of supported devices as the most likely limiting factor.

WI-FI AND PASSPOINT

Amid the extended debate about whether DAS, small cells or CBRS offer the best path forward for providing on-campus and in-building cellular support, most analysts conclude that all these approaches have significant flaws. Hence there is renewed interest in a fourth solution: using existing Wi-Fi WLANs and relying on the dual-radio (cellular & Wi-Fi) features of modern cellular devices.

Most private organizations have a well-established enterprise-grade WLAN: over the last decade, theses have become denser, with closely-spaced access points supporting high rates. Even if an older WLAN has coverage gaps, or access points are due an upgrade, these are well understood issues. Engineering expertise is readily available to bring any enterprise WLAN up to modern standards, at a reasonable cost.

Meanwhile, cellular operators have been supporting data offload to Wi-Fi for many years. An obstacle exists in terms of seamless roaming: the solution will not be widely-successful if the user has to configure access to Wi-Fi every time a new building is entered. However, automatic, silent connection to Wi-Fi is now available through a Wi-Fi Alliance feature named Passpoint.

Passpoint was first released in 2012, and cellular operators have been steadily increasing support over the years. Aruba's recent trials have shown that more than 50% of US cellphones are already configured for Passpoint, and the adoption curve is rising sharply.

Seamless Cellular Handoff

The Passpoint architecture covers three areas. Access points are configured to advertise details of reachable cellular operators. Devices are programmed (usually by the operator) to recognize these advertisements, and to authenticate automatically where they are available. And the WLAN must be connected back to the operators' core networks to carry authentication traffic.

In operation, a Passpoint-configured phone will recognize a Passpoint-capable access point, and initiate authentication. The access point forwards the authentication traffic to the operator's server, and the device is authenticated using its SIM credentials. When the access point is authorized (part of the 802.1X security protocol) it connects the phone to the Internet. The Wi-Fi connection is secure and multiple operators can be supported on a single SSID with no additional equipment.

So far, we have described a 'data offload to Wi-Fi using Passpoint' scenario. The second feature that makes Passpoint a viable solution for on-campus cellular coverage is Wi-Fi Calling. This is a standard feature, now supported by most cellular operators world-wide, where voice calls and text messages can be received and sent over a Wi-Fi connection.

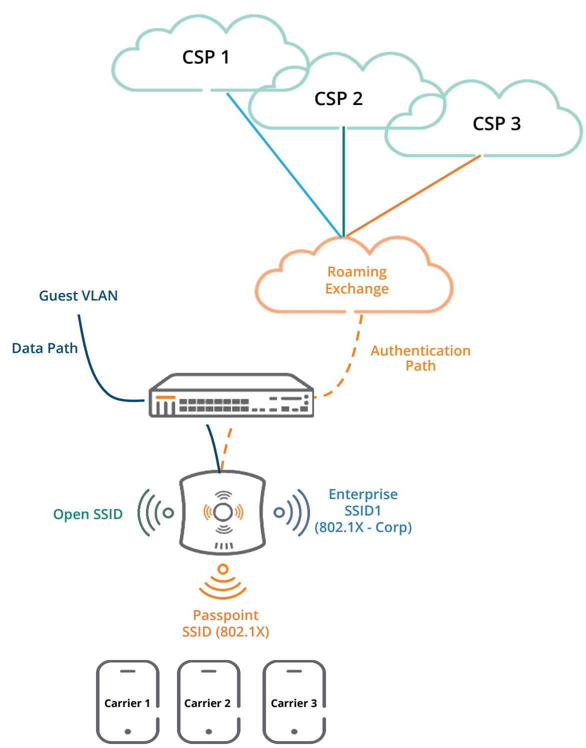

Figure 1: Example of a Passpoint

The combination of Wi-Fi Calling and Passpoint offers a comprehensive solution to on-campus coverage for cellphones and other consumer devices. It can be enabled over most existing enterprise WLANs with a software upgrade and a RADIUS authentication link to the participating operators.

Passpoint still has some challenges. While the majority of consumer devices are ready-configured to roam with Passpoint, some are not. And some features like emergency caller location are not on parity with the cellular network. Nevertheless, we believe it represents a comprehensive solution for enterprise customers, without the many drawbacks of the cellular-technology alternatives.

A MENU OF INCOMPLETE SOLUTIONS

In conclusion, private organizations wishing to provide better coverage for cellular devices coming on-campus can choose from a number of solutions - but all have drawbacks.

The only established solution today is DAS. While DAS is indisputably effective, there are two key challenges:

1) DAS is expensive: A single-carrier solution will cost more than a state-of-the-art WLAN covering the same area, and the cost escalates with the number of carriers supported, so a four-carrier solution will be 4x or more the cost of the WLAN.

2) DAS is inflexible: Upgrades to new technology will be cumbersome and expensive. For this reason, it has had limited success in terms of growing its world-wide market.

Meanwhile, small cells offer the promise of flexible in-building coverage similar to Wi-Fi, but the inconvenient realities of mixing licensed-spectrum infrastructure with private organizations throw up many obstacles to adoption. The most significant, again, is the multi-carrier dilemma: can national cellular operators develop cooperation to the extent that they are willing to roam onto other operators' infrastructure on a building-by-building basis? Although the standards exist, we have seen no movement on the commercial plane, these roaming agreements do not exist today at scale in the enterprise. Technical difficulties in coordinating frequency allocation with the overlapping macro network, and supporting internal and micro-macro handovers are still tricky problems for the small cell.

CBRS is a relatively new development, and on its face offers a new opportunity for private organizations to secure dedicated spectrum at a reasonable cost, and to build their own networks. But there are significant fundamental risks associated with the underlying rules for this band. For supporting visitor and guest devices on-campus, it will take many years for the installed base of cellular devices to fully-support CBRS frequencies, and authentication agreements with cellular operators will be required to support consumers roaming into the campus. None of this exists today, so it remains to be seen when CBRS will become a viable solution.

Finally, the Wi-Fi Calling with Passpoint solution is a different approach. Major US operators have been using Passpoint as their preferred roaming solution for some years, so the majority of consumer devices are already able to roam onto a Passpoint network. With Wi-Fi Calling, all cellphone services are supported, with no user configuration required. Passpoint is a well-understood technology that can be enabled on existing enterprise WLANs at very low incremental cost compared to the alternatives described above. Wi-Fi is a world-wide single standard, inherently capable of supporting multiple operators with no additional equipment.

Technology and markets are changing rapidly, and all the solutions described above will be extended in coming years. Organizations seeking to improve on-campus and in-building cellular coverage have difficult choices. They should carefully examine all the alternatives before making a decision.

For more information on Passpoint and Wi-Fi Calling, the Solving the Indoor Cellular Coverage Problem white paper provides additional detail.

aruba
a Hewlett Packard Enterprise company

WP_Solvingtheindoorcoverageproblem_SK_012920 a00090760enw