# The Room With No Walls: Designing Wi-Fi for Stadiums




https://www.linkedin.com/pulse/room-walls-designing-wi-fi-stadiums-jarryd-de-oliveira-hrkye/

## Every other building gives you walls to separate your cells. A stadium gives you an open bowl and sixty thousand phones.

In every environment we have looked at so far, the building has quietly helped you.

Classroom walls separate one cell from the next. Hotel rooms isolate their guests. Hospital construction, for all its difficulty, breaks the space into pieces. Even warehouse racking creates aisles.

A stadium takes all of that away.

It is an open bowl. Tens of thousands of people sit in one continuous space with very little physical separation between them, and every access point has the potential to hear a large number of other access points using the same channel.

That is the constraint that decides the design.

In a stadium there are no walls to separate your cells.

So you have to build them yourself.

# Coverage is free. Capacity is everything.

Here is the thought experiment that explains why stadium design is different.

Stand in the middle of the pitch with a powerful access point and a high-gain antenna, turn it up to full power, and you could probably put a usable signal into most of the seats.

Every heat map would be green.

And the network would fail completely the moment the gates opened.

Signal strength was never the problem. One radio, or even a handful, cannot share airtime between tens of thousands of devices trying to transmit at the same time.

The coverage exists.

The capacity does not.

In most buildings, coverage is the thing you design for and capacity is the thing you check.

In a stadium it is the other way around.

Coverage often comes relatively easily because nearly every seat can be placed within range of something.

Capacity is the whole design.

# Manufacturing the walls

If the building will not separate your cells, the RF design has to.

The objective is to carve one enormous open space into a large number of small, controlled cells, each serving a manageable number of seats and each using spectrum in a way that allows aggressive reuse without creating unnecessary contention.

The tools for doing that are familiar.

They are simply used far more deliberately than in most other environments.

**Directional antennas.** Specialist directional antennas with carefully selected beamwidth and gain that place energy into a defined block of seats while limiting how much spills into neighbouring sections. In an open bowl, the antenna pattern becomes the wall.

**Low transmit power.** Turning power down helps shrink each cell so that it covers only the seats it is intended to serve. More radios at lower power, rather than fewer radios shouting across the venue.

**A disciplined channel plan.** Narrow channels, planned deliberately and usually held static, because the goal is reuse across a huge number of cells rather than headline speed for any individual client. Allowing an automated algorithm to reshuffle channels during a live event can introduce risk into an environment where the reuse plan has been carefully engineered.

**Controlled overlap.** Enough overlap to maintain usable mobility between cells, but no more than necessary. Excessive overlap in an open bowl quickly becomes co-channel contention across entire seating blocks.

Put simply, you are using physics to recreate the walls the architect left out.

# Where the radios go

In a stadium, mounting position is one of the biggest design decisions you will make because it determines how effectively you can contain each cell.

There are three broad approaches, and large venues often use more than one.

**Overhead.** Access points mounted on roof structures, gantries or catwalks, aimed down into the seating with directional antennas. This keeps equipment out of reach and can work very well, but the distance to the client is large and the angle matters, so antenna selection, mounting height and aiming become critical.

**Under the seat.** Access points mounted beneath the seating, radiating into the crowd above. The radio is very close to the clients it serves, and the people sitting around it provide a significant amount of additional RF attenuation and containment.

**In the handrail.** Access points integrated into handrails, step risers or similar structures, often used where seating geometry, aesthetics or access limitations make other approaches difficult.

Each comes with practical constraints that have nothing to do with RF.

Seats get stood on.

Drinks get spilled.

Seating areas are jet-washed after events.

Equipment has to survive weather in open-air venues, remain accessible for maintenance and, in many cases, stay visually unobtrusive so it does not spoil the appearance of the venue.

Colour matching, low-profile enclosures and hidden mounting are real design requirements here, not afterthoughts.

# The crowd is the building

This is the part that surprises people who have not worked in venues before.

A human body is mostly water, and water absorbs RF energy.

A sold-out stadium contains tens of thousands of absorbers packed closely together, and that additional attenuation is a large part of why high-density designs such as under-seat deployments can achieve useful cell containment.

It also means an empty stadium is not the stadium you are designing for.

Survey an empty bowl and RF will often travel significantly further than it does when the venue is full. Cells appear larger, neighbouring sections can interact more strongly and the overall propagation environment looks very different.

Then fill the same bowl with people and everything changes.

The crowd introduces substantial attenuation, but it also introduces tens of thousands of active client radios, all contending for airtime and generating traffic.

So the occupied state is not simply a quieter RF environment.

It is a completely different one.

If that sounds familiar, it should.

It is the warehouse problem again, with people in place of pallets.

The occupied state is the one that matters, and it is the hardest one to measure.

# Everyone does the same thing at the same moment

Stadium traffic is entirely event driven.

Large numbers of devices sit relatively quietly for long periods, then act together at predictable moments.

Gates opening.

Kick-off.

A goal.

Half-time.

A controversial decision.

The final whistle.

And the traffic profile at those moments is unusual.

Most enterprise networks are built around users who consume far more data than they send.

A stadium crowd at a big moment can behave very differently.

Thousands of people capture the same moment and try to upload it at almost exactly the same time, posting photos, videos and messages before the person sitting next to them does.

Meanwhile, the people are there to watch the event, so sustained internet video streaming may be less dominant than it would be in many other environments.

Capacity planning therefore has to account for sudden upstream demand, not just average downstream consumption.

It also has to account for everything else sharing the air.

The public guest network may carry most of the clients, but the same venue is also running point-of-sale terminals, ticket scanning at the turnstiles, staff communications, broadcast and media operations, security systems and sometimes wireless video.

Those operational systems cannot be allowed to suffer because the crowd decided to post the same goal at once.

Segment them properly.

Give them their own networks, their own policies and appropriate priority, and make sure the design protects them from guest load rather than forcing them to compete with it.

# Plan the capacity honestly

Nobody designs a stadium on the assumption that every device will be transmitting at full rate at the same moment.

The numbers would be absurd.

Stadium design works on oversubscription: an informed assumption about what proportion of attached devices will be actively demanding bandwidth during peak periods, and how much throughput those active clients are likely to require.

That ratio depends on the venue, the event, the audience and the services being offered.

The important part is being explicit about the assumptions behind it.

Then protect the shared resources.

Per-client rate limits are common in venues, and their main purpose is not to magically create more airtime.

Their purpose is to stop a small number of heavy users consuming a disproportionate share of the venue's internet or upstream capacity, which can become a bottleneck independently of the RF.

Push capable clients towards the higher bands where there is far more spectrum available.

Use 5 GHz and, where the client population supports it, 6 GHz for capacity, and keep 2.4 GHz limited to what genuinely needs it.

Features that steer clients between bands or distribute them across access points can help, but they are implemented differently by different vendors and behave differently with different client devices.

Treat them as tools to be tested, not settings to be trusted.

# The bowl is only half the venue

It is easy to focus entirely on the seating bowl and forget that a stadium contains several other environments.

Concourses are corridors of moving people that can fill almost instantly and empty just as quickly.

Food, beverage and retail areas carry the point-of-sale terminals that generate revenue.

Hospitality suites and executive boxes have their own expectations, often closer to a premium hotel than a public network.

Press areas and broadcast positions need reliable, protected connectivity.

Back-of-house areas, offices, changing rooms, car parks and external plazas all need consideration in their own right.

Each is a separate design with its own requirements, even though they may share the same controllers, switching infrastructure and spectrum.

# A war story from half-time

The moment I always think of was not in the seating bowl at all.

It was the concourse at half-time.

The bowl design was solid.

The seating cells were well contained and performed well throughout the first half.

Then the whistle went.

Within minutes, a large part of the crowd left their seats together and poured into the concourses, queuing for food, drink and toilets.

The concourse access points had been designed around people moving through the area.

For fifteen minutes they were serving a crowd several times larger, all standing still, all reaching for their phones at once.

The guest network slowed.

More importantly, the point-of-sale terminals serving those same queues were sharing the same RF environment, and card payments started taking far longer than they should.

The queues grew.

And every slow transaction meant lost revenue at exactly the moment the venue made most of its money.

The fix was to design the concourses for their true peak rather than their average, give payment traffic its own protected network with appropriate policy and priority, and validate the design during a real half-time rather than a quiet walkthrough.

The lesson was simple.

In a stadium, the crowd moves as one.

Your design needs to follow it.

# The deadline does not move

Stadium projects have a constraint many wireless projects do not.

The date is fixed.

The opening fixture, the start of the season or the major event is going to happen whether the network is ready or not.

There is no extension.

That affects everything.

Some design options you would prefer may simply not be possible within the available time.

Installation windows fall between events.

Access to parts of the venue may be limited.

Building and safety requirements may dictate low-smoke cabling, restrict certain mounting locations or require specific approval processes.

A realistic stadium design is not just one that works on paper.

It is one that can actually be installed, tested and tuned before the gates open.

# Validate with a full house

This is where stadium validation differs from almost everywhere else.

An empty survey can confirm coverage, check antenna alignment and establish a baseline.

It cannot tell you exactly how the network will behave when the seats are full because both the RF environment and the traffic load change when the crowd arrives.

So plan to validate under real conditions.

Monitor the first live events closely.

Watch channel utilisation, retries, airtime utilisation, client counts and throughput section by section.

Pay particular attention to the moments that matter:

Gates opening.

Kick-off.

Half-time.

The final whistle.

Measure at seat height rather than simply walking the aisles.

And expect to tune.

In a venue like this, the first sold-out event is not the end of the project.

It is the first real test of the design.

# Final Thoughts

Every environment in this series has had a constraint that decides the design.

In a stadium, that constraint is the absence of walls.

The building will not separate your cells, so you have to separate them yourself with antennas, transmit power, mounting position, channel planning and, surprisingly often, the crowd itself.

So design for capacity, because coverage will usually be the easy part.

Carve the bowl into small, controlled cells.

Choose mounting positions that allow you to contain them.

Plan for upload surges at the moments that matter, not average daily demand.

Protect the operational systems from the guest network.

Design the concourses for the rush, not the stroll.

And never trust an empty stadium to tell you how the full one will behave.

Because in this environment, sixty thousand people are part of the RF design whether you planned for them or not.

---

*This is part of a series on designing wireless for real environments, where the constraint that dominates the building decides the design.*

*Jarryd De Oliveira, CWNE #594*

---

Revision #2
Created 24 September 2026 04:23:45 by Jarryd
Updated 25 September 2026 04:58:18 by Jarryd