Good Wi-Fi Isn't an Accident: What It Actually Takes to Design a Wireless Network

https://www.linkedin.com/pulse/good-wi-fi-isnt-accident-what-actually-takes-design-de-oliveira-rh0de/

Most bad Wi-Fi was never designed. It was installed.

Someone bought access points, mounted them where the cabling was convenient, left everything on default settings, and hoped. Then the complaints started. Dropped connections, dead zones, devices that won't roam, and a network team stuck firefighting something that was broken before it went live.

Designing a wireless network that actually works isn't about buying the latest hardware. It comes down to three things: understanding how RF behaves, understanding what the business actually needs, and using the right process to bring the two together. Get those right and the vendor choice almost takes care of itself.

1. Understand the Physics First

Wi-Fi lives or dies on radio frequency behaviour, and RF doesn't care about your project deadline.

The bands are not interchangeable. 2.4 GHz gives you reach but only three usable channels and a band crowded with everything from legacy devices to non-Wi-Fi interference. 5 GHz provides significantly more spectrum and cleaner air, at the cost of shorter range. 6 GHz introduces a huge amount of new spectrum with no legacy clients, but the coverage footprint shrinks again. Every band represents a trade-off between coverage and capacity, and a good design uses each one deliberately rather than treating them as one large pool of spectrum.

Signal drops fast, then slowly. Free space path loss means you lose the most signal in the first few metres, then approximately 6 dB every time the distance doubles. Higher frequencies start with a larger initial loss, which is exactly why a 6 GHz cell is smaller than a 2.4 GHz one from the same access point.

Coverage and capacity are different problems. A warehouse may require large coverage cells, while a conference space may need many smaller cells to support hundreds of simultaneous users. Designing for one without considering the other is a common reason wireless networks struggle under load.

Signal strength alone tells you nothing. What matters is signal-to-noise ratio. A -65 dBm signal in a quiet RF environment performs well. The same signal next to a noisy production line or surrounded by wireless cameras may be unusable. Coverage, SNR, and co-channel interference together determine whether a design works, not a single colour on a heatmap.

If a design conversation never mentions attenuation, transmit power, capacity planning, or channel reuse, it isn't a design conversation. It's a shopping trip.

2. Requirements Before Radios

This is the step most organisations skip, and it's where CWDP methodology gets it right: design starts with requirements, not equipment.

Business requirements come from stakeholders. What is the Wi-Fi for? Barcode scanners on a warehouse floor, voice handsets in a hospital, laptops in an office, and AMRs in a distribution centre all need very different networks. "Everywhere, fast" is not a requirement.

Technical requirements define the detail. Which client devices matter most? Design for the worst device that has to work, not the best one. A modern laptop will cope almost anywhere. A handheld scanner with a weak radio, used inside a metal rack aisle, will not. That device sets your coverage targets, not your phone.

Environmental reality shapes everything else. Brick, concrete, glass, metal racking, changing stock levels, machinery generating RF noise, and building layout all influence wireless performance. These factors must be measured or modelled, not guessed.

The output of this stage is a set of measurable RF targets: minimum signal level, minimum SNR, secondary coverage for roaming, and acceptable interference limits, defined per band and per area. If you can't write those numbers down, you can't validate the network later, and you'll have no way to prove it's working as designed.

3. Design, Validate, Then Trust the Data

With physics understood and requirements defined, the process is straightforward, and deliberately boring.

Build a predictive model using accurate floor plans, real wall materials, and realistic AP mounting positions. The model is only as good as the information going into it. A plan drawn on a guessed scale with default wall types will produce a confident, wrong answer.

Then validate against reality. Survey the site properly, measure attenuation through the walls you care about rather than assuming it, and confirm coverage at the height and locations where devices will actually be used. An AP tested at two metres and installed at ten behaves like a different network.

After deployment, perform a validation survey. This is where design assumptions meet reality. Coverage, SNR, roaming behaviour, interference levels, and application performance should all be measured against the original requirements. If you don't validate, you're guessing.

Finally, configure with intent. Transmit power matched to your weakest client, sensible minimum data rates, channel widths the environment can support, and a channel plan that minimises contention. Defaults are a starting point for the vendor, not for you.

Final Thought

None of this is vendor magic, and none of it is optional.

The organisations with reliable wireless aren't the ones that spent the most money. They're the ones that treated Wi-Fi as engineering: physics first, requirements second, validation always.

Good Wi-Fi isn't an accident. It's the result of a repeatable design process backed by measurable outcomes.

If your network was installed rather than designed, you already know which one you've got.

Jarryd De Oliveira, CWNE #594.

Revision #1
Created 12 June 2026 04:31:42 by Jarryd
Updated 12 June 2026 04:42:59 by Jarryd