The Network That Cannot Go Down: Designing Wi-Fi for Hospitals



https://www.linkedin.com/pulse/network-cannot-go-down-designing-wi-fi-hospitals-jarryd-de-oliveira-jzgee/?trackingId=jvabS9YFT7KivFtW5MoylA%3D%3D

In most buildings, bad Wi-Fi is an inconvenience. Here it sits on a clinical risk register.

In a school, a bad design produces complaints.

In a hotel, it produces refunds.

In a hospital, it produces incident reports.

That single difference changes every decision you make, because the wireless network is no longer just carrying email and video calls. It is carrying infusion pump data, vital signs telemetry, nurse call, bedside observations, barcode medication administration, clinical images and voice handsets that staff use to reach each other quickly.

None of that can be treated as an acceptable failure.

And here is the part that makes hospital design genuinely different from almost anything else in this series.

You cannot simply take the network down to fix it, and you cannot necessarily fix the clients.

The client estate spans thirty years, and you are not allowed to change it

Everywhere else, the answer to an awkward legacy device is eventually to replace it.

In a hospital, that device may be a regulated medical device.

An infusion pump, telemetry transmitter or diagnostic cart may have been validated by its manufacturer against a specific wireless configuration. Changing what it connects to is not a matter of preference. It may require the manufacturer's involvement, revalidation or a documented clinical risk assessment before anything moves.

So you may be designing for a client estate that includes a brand-new tablet issued last month and a pump that expects 2.4 GHz, an older security method and possibly low data rates, with years of service life still left in it.

The instinct built from almost every other environment is to modernise that away.

Disable 2.4 GHz.

Raise the minimum data rates.

Move everything to WPA3 and 6 GHz.

In a hospital, applying that instinct without checking first is how you take clinical equipment offline.

The answer is not to drag the entire estate down to the capability of the oldest device, and it is not to force the oldest device to keep up.

It is to separate them.

A dedicated SSID and VLAN for legacy clinical devices, with the settings those devices actually require and tightly controlled policy around them, allows the rest of the hospital to move forward while the pumps keep working.

Then you find out what is actually on the network before changing anything.

Not what the asset register says.

What is associated right now, by device type, radio capability and security method.

That inventory should be one of the first deliverables in a hospital wireless design, and it is often the thing that changes the plan.

Roaming is a clinical function here

In an office, a bad roam interrupts a call.

Here, the things that roam may be a nurse carrying a voice handset down a corridor, a workstation on wheels being pushed from bay to bay during a drug round, or telemetry attached to a patient who is walking between departments.

Those sessions have to survive the transition consistently, in a building made up of long corridors, fire doors, lift lobbies, stairwells and varying construction materials.

That means coverage overlap, cell sizing and roaming-assist features all matter more than usual.

More importantly, they need to be proven with the actual clinical devices rather than a laptop.

The handset your staff carry has its own roaming logic, and it was almost certainly not written with your building in mind.

Then there is location.

Many hospitals rely on real-time location services for equipment tracking, staff safety badges and infant protection.

Location is not the same requirement as coverage.

It depends on how many access points can hear a tag, the quality of those observations and the geometry of where those APs sit. That sometimes means placing radios where a pure coverage design would not bother.

If location is in scope, it has to be part of the design from the beginning.

Retrofitting location accuracy into a coverage-only deployment is expensive and rarely satisfying.

The building is actively hostile

Hospitals contain construction you will not find in most other environments.

Radiology rooms may be lead lined.

Theatre suites can have dense construction and sealed finishes.

There are plant rooms, fire compartments, thick structural walls and long service corridors that can behave like waveguides.

And then there is the MRI suite.

An MRI room is a shielded enclosure, effectively a Faraday cage, built to keep RF interference out.

Your signal stops at the boundary, and that is by design.

Coverage inside becomes a specialist conversation involving the scanner manufacturer, medical engineering and the people responsible for that environment. It is not something you solve by simply adding another access point nearby.

You also cannot put equipment wherever you like near the magnet.

Ferrous material and strong static magnetic fields do not mix, and the safety zoning around an MRI scanner is not negotiable.

So the wireless design has to accept hard edges.

There are places in a hospital where wireless coverage may stop, and sometimes the correct engineering answer is to document that boundary clearly and design the clinical workflow around it rather than promising coverage that cannot safely or realistically be delivered.

Infection control decides where things go

In most buildings, mounting is an engineering question.

In clinical areas, it can become a governance question.

Clean areas may not permit exposed containment or surface-mounted equipment.

Ceilings in some spaces are sealed and are not opened casually.

Cleaning regimes may involve chemicals that punish the wrong enclosure.

Older estates may also contain asbestos, turning what looks like a simple cable route into a managed process.

Access is its own constraint.

You may need a permit to work, an escort and an agreed slot that does not disturb patients or clinical activity.

Some areas can only be entered when they are not in use, which may mean three in the morning or during a planned closure months away.

None of this changes the RF.

All of it changes what is achievable, and when.

That is why the estates and infection-control conversations belong at the beginning of a hospital wireless design rather than at installation.

The people who tell you where you cannot mount an access point are doing you a favour early and ruining your programme late.

There is no easy maintenance window

Most other environments have a quieter period.

A school has holidays.

A hotel has shoulder seasons and empty floors.

A hospital runs continuously.

There is no week when the wards are conveniently empty, and a controller upgrade is not simply a technical decision taken by a network team on a Tuesday evening.

It may need clinical approval, a documented rollback plan, change advisory board approval and agreement from the teams responsible for the systems riding on top of the WLAN.

That reality has to be designed in rather than discovered afterwards.

Build for resilience so that a single failure does not remove service from a clinical area.

Think about what happens when a switch fails.

An uplink.

A controller.

A WAN or data-centre link.

Ask whether the wireless service continues to support patient care while that failure is being repaired.

Plan changes in small, reversible steps.

Group related changes where appropriate.

Get them approved as part of a controlled programme rather than repeatedly asking for emergency windows because dependencies were discovered late.

And document what each change touches clinically.

Because the person approving the change is thinking about patients, not packets.

And they are entirely right to.

You will never get a perfect survey

Here is one of the practical constraints that catches people who are used to surveying offices.

You cannot walk everywhere.

Occupied bays, isolation rooms, theatres in use, paediatric and maternity areas, mental health units and anywhere a patient has a reasonable expectation of privacy are not places you simply wander through with a survey rig on a trolley.

So a hospital survey becomes a negotiation.

You go where you are allowed.

When you are allowed.

Often escorted.

And you build the picture from the spaces you can physically reach combined with careful analysis of the construction and layout in the spaces you cannot.

That means accepting that some data will be inferred.

It also means being completely honest in the report about which measurements are observed and which conclusions have been inferred.

Most importantly, validation needs to use the clinical devices that matter.

A voice handset on a real call.

A workstation on wheels pushed along the route a drug round actually takes.

A telemetry device following the journey a patient makes from a ward to imaging and back again.

Those tests tell you whether the design works.

Heat maps tell you whether the coverage exists.

They are not the same statement.

A war story about an optimisation

The example I keep coming back to was not a failed deployment.

It was a tidy-up.

A hospital had a well-intentioned optimisation carried out.

Disable the legacy data rates.

Reduce 2.4 GHz.

Clean up the channel plan.

All defensible practice in a modern enterprise environment.

All recommended in good faith.

Within hours, a set of clinical devices had dropped off the network.

They were older.

They were 2.4 GHz.

They depended on the rates that had just been removed.

And nobody involved in the change had a complete picture of what was still relying on them.

The RF was measurably better afterwards.

The clinical service was worse.

That is the lesson.

Optimisation is not wrong.

Modernisation is not wrong.

But in a hospital, best practice borrowed from another environment is a hypothesis, not an instruction.

Every change has to be checked against what is actually associated to the network, and the estate has to be segmented so that modernising one part cannot quietly break another.

Final Thoughts

Hospital wireless is not harder simply because the buildings are complicated, although they are.

It is harder because the wireless network has become part of a clinical system.

It carries services that can affect patient care.

It supports equipment you may not be permitted to reconfigure.

It operates in a building where you cannot mount freely, survey completely or casually schedule downtime.

And it never really stops.

So start with the inventory, not the heat map.

Segment the legacy estate rather than designing the whole network down to it or forcing devices to modernise beyond their capability.

Treat roaming and location as clinical requirements where they matter, and prove them with the actual devices and workflows.

Accept the hard edges the building gives you and document them honestly.

Design resilience around failures before they happen.

And build a change process that includes the people who are thinking about patients as well as the people who are thinking about packets.

Before any optimisation, ask the question that matters most in this environment.

What is still relying on the thing I am about to remove?

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 14 September 2026 07:12:08 by Jarryd
Updated 14 September 2026 07:15:51 by Jarryd