



UPGROWTH HEALTHCARE PLAYBOOK 05

The Multi-Location Network

How to run acquisition for centres, branches and clinics where the market is a catchment and not a city.

HEALTHCARE

Why this playbook exists

A network of 12 centres is not 1 market with 12 addresses. It is 12 markets, each about 20 minutes wide, and in most networks nobody has drawn a single one of them. The media plan is built on the city, the budget is set at brand level, and the numbers that would show the mistake sit in an operations system marketing never opens.

2 facts decide almost everything about acquisition for a network, and neither lives in a marketing tool. The first is the catchment: how far a patient will actually travel for this service, at the hour they travel for it. The second is capacity: how many chairs, machine hours, consulting slots or theatre days that location has, by shift. Get those wrong and no amount of creative or bidding repairs the outcome, because you are buying demand you cannot serve, in places that were never yours.

What follows are the decisions that fall out of those 2 facts. Where the market starts and stops. Where to stop spending, and why that recommendation is almost never made. Whether the next site adds demand or splits it. Which asset a patient in a hurry actually looks at. And what a report has to look like to survive 12 locations.

Who this is for

Marketing and growth leads at hospital groups, clinic chains, diagnostics networks, dialysis and day-care chains and multi-site clinician practices in India. It assumes you run more than 1 site and that your reporting rolls them up. No particular CRM, no mapping software and no new budget is assumed. Every worked figure here is illustrative and labelled as such.

WHAT IS IN THIS PLAYBOOK

01	The catchment is the market unit	Travel time at the hour the service runs, and the pin code map nobody draws
02	Capacity is the binding constraint	Chairs, shifts and machine hours, and the location whose right budget is 0
03	Cannibalisation, before the lease	The overlap arithmetic, and the split that stays invisible for 3 quarters
04	The listing is the campaign	Map profiles, hours, tracking numbers, and review velocity as a metric
05	A page per location	Why 1 Services page cannot rank for a locality, and the duplicate trap
06	Reporting that survives a network	Never blend, routing as a field, and the 60 day build

01 The catchment is the market unit

Ask a network where it competes and you get a city. Ask where a patient comes from and you get a journey. Only the second can be bought media against.

Define it in minutes, not kilometres

A radius is the wrong shape and the wrong unit. A 6 km ring crosses a river, a rail line and a flyover that runs one way in the morning, and on the ground those are walls. So use travel time, and **compute it at the hour the service actually runs**. A dialysis first shift starts before the traffic does; an evening OPD sits in the worst hour of the day. Draw **3 rings**, at 15, 25 and 40 minutes, separately per service line, then **add the barrier layer by hand**. An isochrone that ignores a level crossing is fiction.

Service line	Journey pattern, and what decides it	Practical catchment
Dialysis and day-care	Several times a week, indefinitely. A slot on a reachable shift	Tightest. Minutes, not kilometres
Pathology and imaging	Episodic, often a fasting morning. Turnaround and home collection	Very tight, and shrinking
ENT and audiology, acute	1 or 2 visits, wanted today. Earliest available	Decided in the map pack
Orthopaedic joint work	Few visits, 1 large event. Surgeon and institution	Wide. Proximity is a tiebreak
Oncology	Many visits, 1 decision. Institution and clinician	Widest, often inter-city

The rule underneath: the more journeys a patient makes, the smaller the market and the more of your advantage is geographic rather than clinical. The collection-network version is on our [diagnostics and pathology lab](#) page.

Now plot your own patients on it

Every registration record carries a residential pin code. Networks hold that data, for years, and almost none have plotted it. Export 12 months with 4 fields, pin code, location attended, service line and month, aggregate to pin code level, and lay the counts over the rings.

3 things come out. **The real catchment shape**, never a circle: a lobe pulled along whichever road moves, typically half the area you targeted. **The gap pin codes**, inside the ring and sending you nobody. That is your addressable demand, and it usually sits in a different direction from the one the media plan assumed, because radius targeting spends its budget where you already win. **The outliers**, outside the ring and sending patients anyway, each with a cause worth finding.

Keep the map on the right side of the data line

Aggregate inside the clinical system and hand over counts. This stops being a marketing artefact the moment a row is a person. The Digital Personal Data Protection Act, 2023 is being brought into force in stages, so check the current notification.

02 Capacity is the binding constraint

Media planning is built on reach. Healthcare delivery is built on units of capacity fixed for the length of a lease, and the collision is settled in favour of the media plan, because that is the one with a monthly meeting. The binding unit is never beds or square feet. It is whatever runs out first.

Service line	Unit of capacity, and how it runs out	What an idle unit still costs
Dialysis	Chairs times shifts a day. Exhausted shift by shift	Rent, rostered nurses, machine lease
Imaging	Machine hours by modality. MRI full while X-ray sits idle	Equipment lease, technologist on shift
OPD and consulting	Slots per consultant per session	Session fee and a room
Day surgery	Theatre days and list length. Surgeon and anaesthetist diary	Theatre staff on shift

Capacity is not 1 number for a location. It is a grid: location, times shift, times modality or consultant. A centre at 92% can be full on both morning shifts and at 55% in the afternoon. It does not need more enquiries, it needs enquiries from people who can come at 2pm, which is a rostering problem. More in [Pain in the Kidney](#) and on our [dialysis and nephrology chain](#) page.

The location whose correct budget is 0

Money pointed at a full location buys enquiries you decline, while cost per enquiry looks normal. You are paying to build a waitlist for a competitor, because somebody told to wait 11 days calls the next result. Meanwhile the location across town with an empty afternoon accrues rent, a roster and a lease whether the slot is used or not. So for a site at capacity on every shift it can staff, the correct media budget is 0 and the money should move across town. That is almost never recommended.

A percentage of media fee cannot produce this recommendation

If your agency is paid a percentage of what it spends, the most valuable recommendation available to a multi-location network, **cut this location to 0**, costs the agency money to make. Nobody has to be dishonest for this to bite: the analysis never gets commissioned, or arrives framed as a reason to add budget elsewhere. It runs internally too, and a team on enquiry volume will not propose a cut either. The test: read 12 months of recommendations and count how many proposed spending less anywhere. If the answer is 0 across a network whose utilisation varied, you are reading an incentive.

- 1. Pay for scope, not for spend.** A fixed retainer against defined work, so a reallocation or a cut is fee-neutral to whoever recommends it. That is the fix. The rest is enforcement.
- 2. Put utilisation in the marketing pack,** from operations, on the same row as cost per patient.
- 3. Set a spend ceiling per location, owned by operations.** Above a stated utilisation threshold it falls automatically and nobody debates it in the room.
- 4. Report reallocation as a line item,** in rupees moved away from constrained sites this month.

03 Cannibalisation, before the lease

Networks rarely expand from a map. They expand from an opportunity: a property comes free, a consultant brings a practice in, a competitor exits. Open a second site inside an overlapping catchment and you have not added a market. You have split 1 demand pool across 2 rent bills.

Why the split stays invisible for 3 quarters

Network enquiry volume: unchanged. Cost per enquiry: unchanged. Total patients: up slightly, which reads as the new site working. Blended cost per patient acquired: unchanged. Every headline number in a marketing pack is structurally incapable of detecting a demand split, because they are sums across locations and a split conserves the sum. The 2 numbers that move are utilisation at the original site and patients by residential pin code, and neither is in the pack unless somebody put it there.

The arithmetic to run before you sign

- 1. Overlap share.** Location A has a 25 minute catchment covering 40 pin codes. Proposed location B, 4 km away, covers 30, of which 22 sit inside A. So 73% of B is territory you already serve.
- 2. Exposed volume.** A sees 600 patients a year at 88% utilisation, and 430 live in the 22 shared pin codes. That is the population at risk, and it is a fact in your registration data.
- 3. Switching assumption, stated out loud.** For a high-frequency proximity service assume half move to the closer site: 215 patients. Write it down, then run it again at a pessimistic 70%.
- 4. Incremental yield.** B reaches 8 pin codes A does not. A runs at roughly 15 patients per pin code a year in its core; discount edge pin codes to 60%. About 72 genuinely new patients a year.
- 5. Net it.** A falls from 600 to 385 and utilisation from 88% to about 56%. B opens at 287, of which 215 were already yours. The network goes from 600 to 672, 12% growth, against a second set of fixed costs.

Illustrative arithmetic. Figures invented to show the method, not drawn from any client engagement.

The board paper says 12% patient growth. The P&L says 72 incremental patients are carrying a second rent, roster and licence set, while the original site sits at 56% and looks like a performance problem it does not have. So 1 division decides it: incremental patients times contribution per patient, against incremental annual fixed cost. If that needs 3 years to clear, say so in the paper rather than in month 9.

When cannibalisation is the right answer

Sometimes it is. Defending a corridor before a competitor takes the lease. Releasing capacity at a site that is full and cannot be extended, which only works if unserved demand remains in its catchment to refill the seats, and that is what the gap pin codes tell you. Or cutting travel time to the level a high-frequency service needs. All 3 are legitimate. They have to be the stated reason, with the split in the forecast, so the original site being down 35% is a plan and not a crisis.

If your network already overlaps, differentiate the catchment rather than the creative. Geo-exclude the shared pin codes from the site you want to unload and target them only from the receiving one. Route them by written rule rather than by whoever picks up. And move service lines rather than patients: wide-catchment specialty at 1 site, high-frequency proximity service at the other.

04 The listing is the campaign

For a proximity-decided service the map profile outranks the website. Somebody searches for a blood test near me at 6:40am, or an ENT clinic open now on a Sunday. They get a map, 3 results and a row of buttons, and decide inside that box. So the highest-leverage asset a network owns is its 12 listing profiles, usually set up once by somebody who has since left. More in [Pain in the Ear](#) and on our [ENT and audiology clinic](#) page.

Field	What usually goes wrong	The fix
Primary category	A broad default like clinic	The most specific accurate primary available
Service list	Empty, or the brochure list	What this site delivers, in patient words
Hours, per day	9 to 5 on a 3-shift site	Real hours per shift. Open now is a filter
Holiday exceptions	Never set. Open on a shut day	12 months of holidays, loaded in 1 sitting
Booking link	Missing, or the homepage	A path that books this location
Photos	Uploaded at launch only	A few a month, of this building
Name, address, phone	Different in every directory	1 canonical string, enforced everywhere
Reviews	Counted, never managed	Velocity, with a named asker per site

The tracking number breaks the thing you are trying to measure

Dynamic number insertion and per-campaign call tracking create a second, third and fourth phone number for the same building, and directories scrape them. The location now has 4 numbers across the web and no 2 sources agree, which degrades exactly the name, address and phone consistency proximity results depend on. Keep the verified listing on 1 canonical number, and never let a tracking number be published as a site's primary number where a scraper reaches it. Use call tracking on paid landing pages and extensions, where you control the page. If the listing number must be trackable, use 1 permanent forwarding number everywhere.

Review velocity, not review count

A count is a stock number. Recency is the only signal an urgent searcher has that a place is currently operating and currently good: 40 reviews with the newest 14 months old reads worse than 9 with 3 from this month. So reviews per month per location is an operating metric, in the same pack as utilisation. A site with no review in 3 months has a problem now: nobody is asking, the asker has left, the link is broken, or something happened patients will not write about.

- **Set a floor per location, not a network target**, or the flagship hides 5 silent sites. Ask at the point of a completed interaction, on a link that opens the correct location profile.
- **Respond without clinical detail**. A public reply confirming somebody was a patient is a disclosure. Keep it to process and an offline route with a named person.

05 A page per location

A query anchored to a locality is a different query from one about a service. Somebody searching for an MRI centre in a named suburb is not searching for MRI, and 1 network Services page carries no locality-specific fact worth ranking. Same for whether an assistant names you: asked where to get a service in a locality, a generative system assembles from sources naming the place and the locality together. The entity retrieved is the location.

What each location page has to carry

- **Name, full postal address and canonical phone number** as text, matching the listing exactly, and **its own hours** per day and per shift, including Sundays and holidays.
- **The services available here, and explicitly the ones that are not.** A page implying every service at every site produces a wasted journey and a 1-star review from somebody right to be cross.
- **Site-specific capability facts.** Modalities, chairs, day-care provision, lab turnaround here, parking, lift, ramp, access, and **photographs of this building**. Not stock, not the flagship.
- **The clinicians who actually sit here**, with the days and sessions they sit. The most valuable block on the page, and the one usually replaced by a link to a network-wide directory.
- **Directions the way a local gives them, and the localities this site serves, named.** Where your pin code map stops being a slide and becomes the content that answers the locality query.
- **A booking route that books this location**, on the number the listing carries, plus payment and empanelment facts, and **structured data for this location**: address, geo coordinates, opening hours, telephone, and a sameAs to this site's map profile. 1 record per site, never 1 for the brand.

The duplicate content trap

12 pages generated from 1 template with the town name swapped is not 12 pages. It is 1 page repeated, and search systems treat it that way: 1 version gets indexed and ranks weakly, the rest get filtered, and your own pages compete for the same locality query. The test takes 30 seconds. Cover the town name and read the page. If you cannot tell which location it is, neither can anything else. Most words on each page have to be facts true only of that site.

Then link it properly. A location index 1 hop from the homepage. Both directions between service and site, because most networks build 1 and wonder why location pages stay orphaned. Cross-links between nearby sites, so a full location can pass demand across town. And no page per locality per location.

Want the per-location audit rather than the checklist

We run every site in your network against the listing fields and page requirements above, cross-check against your own pin code map, and hand back a ranked fix list per location.

Book a 30 minute growth call or write to contact@upgrowth.in

06 Reporting that survives a network

A 12-centre chain reporting 1 blended cost per enquiry has a number that describes no centre and moves for reasons nobody can name. It is not a weak metric, it is a concealing one: the sites that work subsidise the sites that do not. So never blend.

What to report	Grain	Why it earns its place
Cost per patient acquired	Location x service line	The only comparable cost number
Utilisation	Location x shift x modality	On the same row as the cost. That is the point
Media spend	Location	So a reallocation reads as a movement
Generated for, fulfilled at	Location pair	Separates demand created from demand absorbed
Reviews received	Location x month	Velocity as an operating metric. 0 is an escalation
Patients by pin code	Location x quarter	Catchment drift, and what new money bought

When an enquiry can be served by more than 1 location, something decides where it lands. If that is not recorded, location-level cost is not comparable, and the symptom is a site that looks expensive purely because it absorbs overflow from a busier one. So record 2 locations on every enquiry, generated-for and fulfilled-at, and where they differ record why: no capacity, shift unavailable, service not offered, patient preference. That converts centre C is expensive into centre C absorbs 30% of centre B's overflow.

Mechanics: [Pain in the Handover](#).

Cost the catchment, not the campaign

Cost per patient at a location includes what a media report leaves out: listing management, the location page, the review programme, the coordinator who answers, and overhead. Our [patient acquisition cost calculator](#) is where to begin.

The 60 day build

- 1. Days 1 to 20, the map and the grid.** Registrations as counts by pin code, location and service line. Travel-time rings at the modal arrival hour, plus the barrier layer. The capacity grid from operations. Both on 1 page per site. That page is the media plan.
- 2. Days 21 to 40, fix what costs nothing.** Every listing against the 8 fields. 1 canonical name, address and phone string per site. 12 months of holiday exceptions. A written rule on tracking numbers. Review velocity with a named asker. Pages for the 3 biggest gaps.
- 3. Days 41 to 60, move the money.** A spend ceiling per location from the grid, cut to 0 where a site is full on every shift it can staff. The routing field. A pack with utilisation beside cost. A scope where a cut is fee-neutral.

Run the same sequence over 90 days if the network is larger than about 15 sites, or if registration data has to be pulled site by site rather than centrally. Nothing in the order changes, only the first phase gets longer, and the first phase is the one worth being slow about. Add the 4 overlap numbers from section 3 to it whenever a new site is under discussion. A media plan can be changed next month. A lease cannot.

If your systems cannot hold these fields

Do not start a migration. A sheet with location, shift, capacity, actual, spend and enquiries, reconciled weekly for a quarter, produces the pack and tells you what to buy afterwards. The 2 fields worth fighting for first are the routing pair and days to the next available slot. Between them they explain most of what a network average is hiding, and the second is a leading indicator: when the wait at a site starts lengthening, that site is filling up before the utilisation report says so.

Expect month 2 to look worse: the average was flattering 3 sites. After 60 days you have a map of where your patients come from, a capacity grid, a per-location pack, and a fee structure that permits the words spend less.

One structural point decides whether any of it changes behaviour. The pack has to be read by operations and marketing in the same room. A network where marketing reports enquiries to one director and operations reports utilisation to another will produce 2 accurate reports and 1 wrong decision every month, because nobody ever sees the pair of numbers together. Fix the forum before you refine the dashboard.

Start with the map

upGrowth works with hospitals, clinic chains, diagnostics networks and clinician groups on patient acquisition, search visibility and AI search presence. If you cannot show where your patients live and how full each shift is, on 1 page per location, that is the first week of work.

Book a 30 minute growth call or write to contact@upgrowth.in



DIAGNOSIS BEFORE PRESCRIPTION

Start with one number, not one campaign.

Most networks have never mapped where their patients live.
That map takes a week and it usually moves the media plan.
We will build the first version from your own data.

Book a 30 minute growth call

contact@upgrowth.in | +91 97663 34689 | upgrowth.in

UPGROWTH DIGITAL LLP, 31 Cloud-9, NIBM Road, Mohammedwadi, Hadapsar, Pune 411028, India

This document is confidential and intended for the named recipient. It may not be reproduced or distributed without prior written consent from UPGROWTH DIGITAL LLP.