upGROWTH

UPGROWTH HEALTHCARE PLAYBOOK 03

Healthcare Generative Engine Optimisation

How to get a hospital, clinic or clinician cited inside the answer, and why ranking on page 1 is no longer the same as being quoted.

HEALTHCARE

UPGROWTH DIGITAL LLP | upgrowth.in | contact@upgrowth.in
First edition, August 2026 | upGrowth Digital LLP

Why this playbook exists

Your page ranks. It has been on page 1 for that query for 2 years and the traffic report says it is working. Ask an assistant the same question and your name is nowhere in the answer. Both things are true at once, and in health they are true more often than in any other category.

Generative engines do not read a results page and summarise it. They assemble a small candidate set, typically a handful of sources rather than a page of 10, then build an answer out of sentences they are willing to repeat. Ranking gets you to the retrieval stage. It does very little at the selection stage, and health is where selection is tightest. Medical questions sit inside the strictest source handling any of these systems apply, which pushes the shortlist toward government health portals, national institutes, professional bodies, academic publishers and indexed literature. Commercial provider domains are discounted harder here than in any other vertical.

That sounds like a wall. It is a map. There is a large class of questions where the institutional sources have nothing to say, because the answer belongs to you: what your review costs, what records you need, how many working days your report takes, what happens on the day, who the named consultant is and what they are registered as. Those questions carry the commercial intent. This playbook is about winning them, and about not spending a budget on the ones you cannot win.

Who this is for
Marketing heads, digital leads and agency teams at hospitals, clinic chains, diagnostics businesses and clinician groups in India. It assumes a site that ranks for something and a team that can publish. It assumes no GEO tooling. Where a mechanism is inferred from observed behaviour rather than documented by the platform, this playbook says so.

WHAT IS IN THIS PLAYBOOK

01 Ranked is not the same as quotable

Rank asks whether a page is worth showing to a person who can judge it. Citation asks whether a sentence is safe to repeat when the system carries the assertion.

What health-tuned source selection favours

Every major system treats medical topics as a high-stakes class where a wrong answer causes real harm. The pool tilts to government portals, national institutes, professional bodies and peer-reviewed literature. Provider domains sit below them.

You are not competing for a position on a list of 10. You are competing for a slot in a shortlist rebuilt for every question. Being 4th on page 1 gets you nothing.

The question class	Where the answer comes from today	Your realistic position
Condition, cause, drug, report terminology	Government portals, national institutes, clinical references	Do not compete. Provider domains are discounted hardest here.
Which specialist, which test	Institutional sources, then large aggregators	Partly winnable. A department and clinician map reads as the practical answer.
Cost, records, what happens on the day	Nobody authoritative. Aggregators, forums, old news	Yours to lose. The largest unclaimed citation surface in Indian healthcare.
Who should I go to in this city	Directories, listicles, review platforms	Hard and slow. Won by entity strength off your own domain.
Is this clinician credible	Your profile, registries, professional listings	Yours, if the entity is consistent. Lost immediately if not.

Your rank report will keep telling you everything is fine
Position tracking measures a surface fewer of your patients see every quarter, and it does not fall when citation share falls. Impressions soften while positions hold, then clicks soften, and only then do sessions drop.

The 3 things that move a provider domain up the shortlist

1. **Verifiability at the level of the sentence.** A statement carrying its own definition, period, source and named accountable person can be repeated without the system taking on risk. An adjective cannot.
2. **Corroboration you do not own.** A fact stated only on your site is a claim. The same fact on your site, a council registry, a professional body listing and a hospital directory is a resolved fact.
3. **Territory nobody institutional occupies.** No national institute will publish what your second opinion costs or how long your report takes. That is where healthcare generative engine optimisation lives.

02 The health question chain

Nobody starts by typing your name. Health research starts with a word: a symptom, a drug printed on a prescription, or a phrase copied off a pathology report. The brand enters late. A provider first becomes relevant at question 3 or 4, and by then the assistant has assembled a shortlist you are inside or outside.

The chain tells you what each asset is for. A page written for stage 1 competes with a national institute and loses on every run. The same effort spent at stage 3 is usually uncontested.

The 5 stages, and what each is worth to you

#	What the person types	What the engine does with it	The asset that serves it
1	A single word, or a phrase copied off a report	Answers from institutional sources. Names no provider.	Nothing. Do not build here.
2	Is this serious, which doctor handles it, what test confirms it	Explains the pathway, starts naming specialties	A department and pathway page in patient words
3	What does it cost, what is included, how long, what do I bring	Struggles. Falls back to aggregators and forums.	Priced, dated process pages. Your highest return asset.
4	Where in this city, who does this, second opinion options	Assembles a shortlist from entity signals	Entity strength, corroborated off your own domain
5	Is this named clinician credible, what are they registered as	Resolves the person against every source it can find	A complete, consistent, dated clinician profile

2 things change the shape of the chain. Speed: an acute ENT or ophthalmology query can run all 5 stages in an afternoon, while an elective joint replacement or a fertility decision takes months. And identity: for a serious diagnosis the searcher is very often not the patient. It is an adult child, in another city, at 11pm, working from a report photographed on a phone. They ask process questions, and those are the ones you can answer.

> **The exercise that settles the argument, and it takes an afternoon**
> Take 1 specialty. Write out the 6 questions a family actually asks, in their words, from the word on the report to a booked appointment. Run all 6 through 2 assistants and read the answers as the person asking, not as a marketer. You will find 1 or 2 stages where an aggregator is answering with facts you could have published. That gap is your content plan.

In oncology the second opinion question is a chain of its own: what a second opinion involves, which records to send, whether the original team is told, what it costs, how long it takes. Almost none of that is published by providers. All of it is asked. The commercial shape of it is in oncology and cancer hospital marketing.

03 Entity construction, not keyword targeting

Keyword targeting assumes the unit of competition is a page. In health generative answers the unit is usually a person or an institution. The system is not asking which page matches a string. It is asking who this is, whether this one is the same as that one, and whether it can say anything about them with confidence.

That is entity resolution, and it fails in dull, fixable ways. A consultant whose name appears in 6 forms across 9 properties is 6 weak entities rather than 1 strong one. A hospital that lists a department under its internal name, never under the words patients use, is invisible to the question that names the specialty.

The 10 fields, and where each has to match

Field	The rule	Where it has to be identical
Name form	1 canonical string. Fix honorific, initials and spelling once, in writing.	Profile, schema, directories, business profile, bylines, video titles
Council registration	Number and issuing council, in plain text on the profile	Profile, schema identifier, registry entries
Qualification	Exactly as awarded, in the awarding body's abbreviation	Profile, schema, directories, speaker bios
Current affiliation	Institution, department and start year. Date or remove former ones.	Profile, schema, business profile, directories, bylines
Sub-specialty	The words a patient uses, not the internal department code	Profile, department page, schema, directory categories
Location	1 address string per site, matching the business profile exactly	Site, schema, business profile, listings, map entries
Physician schema	name, medicalSpecialty, worksFor, memberOf, identifier, sameAs	Every clinician profile page
MedicalOrganization schema	name, department, address, telephone, medicalSpecialty, sameAs	Home page and every location page
Named medical reviewer	A reviewer distinct from the author, with their own registration	Every clinical and procedural page
Last reviewed date	Visible, held apart from publish and last-modified dates	Every clinical and procedural page

sameAs is the field most teams skip and the one doing the resolution work. It lists URLs where the same entity appears elsewhere: a registry entry, a professional body listing, a directory page, an indexed author profile, a verified channel. Each is a chance for an engine to confirm the person on your page and the person in the registry are the same person.

The entity you are building can resign
If your visibility rests on 3 named consultants, your visibility has a notice period. When they leave, citations follow the name and not the building, and the stale affiliation persists in third-party sources long after the exit. Build the institution as an entity in parallel from day 1. Put process, cost, scope and logistics content under the organisation, and opinion, explanation and technique content under the clinician. A departure then costs you a byline instead of a department.

The consistency audit, in 1 afternoon

- List every property carrying a consultant's name that you do not control. Directories, group sites, society pages, conference programmes, channels, previous employers.
- Record the exact name, qualification and affiliation strings on each. 1 sheet, 1 row per property.
- Pick the canonical form, usually the one on the council registration, because that is the one a reader can verify.
- Fix your own properties first, then request corrections outward. A meaningful share will be ignored, so start where the traffic is.

1 constraint to write into the process. The professional conduct code binding your consultants restricts claiming a specialty without the qualification in that branch. It does not restrict stating the qualification held, the registration number or the department. Entity work sits entirely in the second category, which is why it is the safest high-return activity available to a healthcare marketing team.

International patients run the entire chain inside an assistant, because they have no local network to ask and no way to check a name by asking a neighbour. The shortlist is assembled from whatever is publicly resolvable about your institution and your consultants, before anyone at your organisation hears from them. What that means commercially is on our medical tourism and international patient marketing page.

04 What actually gets quoted

A generative engine quotes a sentence it can lift off your page, set beside a sentence from another source, and stand behind. That filter excludes almost all marketing copy, not because the copy is badly written but because there is nothing in it to carry.

The 6 content types that get extracted in health

- **Definitional.** What a second opinion is. What day care means on your invoice. What a package includes and excludes.
- **Process, step by step.** What happens at each stage, in order, with real times and the actual counter. The most quoted format we see, and the least published.
- **Records and preparation.** Which documents a review needs, in what format, and what happens if 1 is missing.
- **Cost and inclusions.** The number, what it covers, what changes it, and the cashless position. The most asked and least answered question in Indian healthcare.
- **Duration and turnaround.** Working days to a report, days to the next slot, hours a scan takes. Bounded, checkable, dull, quoted constantly.
- **Scope and capability.** Tests run in house, equipment by make, accreditation numbers, languages at the desk, consultants on the roster.

Notice what is not on that list. Nothing about outcomes, nothing about efficacy, nothing requiring a clinical claim. The extractable surface in healthcare is almost entirely operational, which is the part your team already knows and has never written down.

A rewrite, so the difference is visible

The paragraph as most hospitals write it

With state of the art technology and a team of highly experienced specialists, our Centre of Excellence delivers world class care with a patient first approach, making us the preferred destination for second opinions across the region.

There is not 1 checkable fact in 37 words. Every load-bearing word is an adjective supplied by the party selling the service. Lift any clause out and it means nothing standing alone, which is the test an extractive system applies. A compliance reviewer would strike the same words for the same reason: nobody outside the building can verify them.

The same paragraph, rewritten

A second opinion here is a paid consultation with a consultant in the relevant specialty. You send the pathology report, the imaging on disc or by upload, the discharge summary and the current prescription list. The consultation is scheduled within 3 working days of complete records being received. The written opinion is issued within 2 working days of the consultation and shared with you, and on request with your referring doctor. The fee is stated on this page and does not change if the opinion agrees with the original plan.

Illustrative rewrite. The timings are placeholders for your own service level, not a benchmark.

- Every sentence carries 1 fact and survives being lifted out of the page on its own.
- The specificity moved out of the claim and into the process, the records and the timing.
- Nothing is asserted about what the opinion will conclude, so there is no clinical claim to defend.
- The last sentence answers what the family is anxious about, which no competitor has published.

Inventing a statistic is now actively counterproductive

The old trick was to publish a confident number because it looked authoritative and nobody would check. Systems increasingly discount unsupported medical assertions, and an unsourced figure on a health page reads as a reliability signal about the whole page rather than that 1 sentence. A fabricated number does not simply fail to help. It drags down the true sentences beside it. A figure with no denominator, period, definition or source is also the easiest thing on your page for a consumer regulator to attack. The page that survives a compliance review and the page that gets quoted are the same page. That is 1 test applied by 2 different readers.

Playbook 02 in this series works that through: the 4 statutes and 1 professional code binding Indian healthcare content, a 30 second test a writer runs against a single sentence, and the medical review log that sits behind a reviewed-by badge. Run that check before you scale content, because publishing more uncitable pages is a faster way to spend the budget. The narrative version is Pain in the Prescription.

5 formatting rules that decide extraction

- Put the question in the heading in the words a patient uses, and answer it in the first 40 to 60 words underneath. Elaborate after, never before.
- 1 idea per sentence. Compound sentences do not survive extraction.
- Use a table when the thing is a matrix and a list when it is a list. Do not write out 6 rows as a paragraph.
- Date the page visibly and name who reviewed it. An undated health page is a weaker candidate than a dated page saying the same thing.
- Answer the negative questions: what is not included, what changes the price, what needs a referral first. Nobody else writes them, so they get quoted.

Catalogue businesses have the largest version of this opportunity. A diagnostics chain with several hundred test pages has several hundred chances to publish scope, preparation, turnaround and price, and most publish a test name and a rate. What that looks like as a programme is on our diagnostics and pathology lab marketing page.

05 Measuring something that passes no referrer

This is the part most GEO material skips. A generative answer usually passes no referrer: the assistant uses a fact from your page and the person never arrives.

A second problem is specific to health and it removes the usual escape route. In classic search a lost organic click can be bought back. In AI answers there is no slot to purchase, and health is a restricted advertising category, so paid inventory is constrained exactly where the stakes are highest. This is a build problem, not a media problem, run as a generative engine optimisation programme rather than an audit.

The 4 things you can actually measure

- **Citation share on a frozen prompt set.** Monthly, same week. A panel, not a market.
- **Branded and named-clinician search.** Discovery you cannot otherwise see. Weekly, on a 12 week roll.
- **Direct sessions to deep pages.** Nobody types a URL for a records checklist. Monthly, by URL.
- **Enquiry source.** Add an AI option to your form this week. Your only first-party evidence.

How to build a prompt set worth running

1. **40 to 60 prompts, then freeze them.** Below 40, 1 erratic answer moves the whole number. Above 60, nobody runs it in month 3. Changing a prompt ends the series.
2. **Cover all 5 stages, weighted to 3, 4 and 5.** Write them the way a patient types: full questions, spelling variants, Hinglish where your patients use it.
3. **Name competitors in 6 to 8 prompts, and include 5 to 8 you expect to lose.** The losing set is your control. Log the answer text, not a yes or no: engine, date, named, linked, primary or supporting, and the sentence attributed to you.

> **Be honest about what this is, including with your board**
> These systems are non-deterministic. The same prompt can return a different answer 10 minutes later, and models change without notice. A single run is noise. A direction across 3 or more runs on a frozen set is a signal. Report it as a panel measurement, never as a rank.

Want the prompt set built rather than described

We build the frozen prompt set for your specialties and cities, take the baseline across the major engines, and hand you the log of who is being quoted.

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

06 The 60 day build

The order matters more than the pace. Entity work compounds, and baseline comes first, because from day 8 you can no longer measure what you started with.

5 blocks, in this order

Days	The work
1 to 7	**Baseline before you touch anything.** Write and freeze the 40 to 60 prompts with a version date and an owner. Run the set across the major engines, logged out, and log answers as text. Record which domains recur most. Add the AI option to your enquiry form.
8 to 20	**Fix the entity.** Run the consistency audit across every property carrying a consultant name, then agree the canonical name, qualification and affiliation strings and publish them on your own site first. Add Physician schema to clinician profiles and MedicalOrganization schema to every location page, sameAs populated on both. Put a visible last-reviewed date and a named medical reviewer into the template.
21 to 40	**Publish the process set.** Pick the 2 specialties with the longest consideration chain and the highest value per patient, not the 2 with the most traffic. For each, publish the 6 extractable types: definitional, process, records, cost, duration and scope. Question in the heading, answer in the first 40 to 60 words. Run every page through the compliance sentence test before it ships.
41 to 50	**Corroborate off your own domain.** Correct the canonical clinician details on the 5 external properties with the most inbound traffic. Populate sameAs: registry entries, professional body listings, indexed author profiles, verified channels. Where a consultant speaks or presents, make the name string and affiliation match your canonical form.
51 to 60	**Run it again and read it properly.** Re-run the frozen set in the same conditions and compare prompt by prompt, never in aggregate. Look for the sentence, not the link: being paraphrased without attribution still means you were selected. Then read the enquiry source field.

What moves first, and what will not move at all

Expect entity and verification questions to shift first, sometimes inside 3 weeks, because you are supplying facts nothing else could supply. Process and cost questions follow over 1 to 2 months. The best hospital in a city question may not move in 2 quarters. Judge the programme at day 60 on the wrong class and it gets cancelled.

Diagnosis before prescription

upGrowth works with hospitals, clinics, diagnostics businesses and clinicians on search visibility and AI search presence. Finding out whether you appear inside the answer takes a week.

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

upGROWTH

DIAGNOSIS BEFORE PRESCRIPTION

Start with one number, not one campaign.

Pick 40 prompts a patient would actually type. Run them once.
That log tells you whether you exist inside the answer.
We will build the prompt set with you in 30 minutes.

Book a 30 minute growth call

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