TEMPLATE · FREE RESOURCE

The Niche Evaluation Scorecard

A weighted scorecard for comparing 3-5 candidate niches on eight criteria, from pain severity to expansion path. The Excel file computes weighted scores, ranks candidates and flags any below the cut-off.

By **Ratomir Jovanovic** · ratomir.com/resources/niche-evaluation-scorecard

---

This scorecard is for founders, CEOs and product leaders who have more than one plausible market to focus on and need to pick one. It works for a new company choosing its first beachhead, for an established product looking for its next segment, and for AI products in particular, where the technology often fits a dozen use cases and the real risk is spreading thin across all of them.

The outcome is a ranked shortlist with a written rationale that your team and board can challenge line by line. You will not get certainty from it. You will get a structured argument, a clear view of where your evidence is weak, and a decision that is easier to revisit when new information arrives. The download includes the PDF guide and the Excel scorecard, which calculates weighted scores, ranks your candidates and flags any that fall below the cut-off.

How to use this

1. Write down 3 to 5 candidate niches. Define each one tightly enough that you could build a list of 50 named companies or people in it. "Healthcare" is not a niche; "independent physiotherapy clinics with 3 to 15 practitioners" is.
2. Review the eight criteria below and adjust the weights in the Excel scorecard if your situation demands it. Agree on weights before anyone scores, so the weights cannot be tuned to favour a preferred answer.
3. Gather evidence for each candidate. Aim for at least five conversations per niche plus desk research. Note the source next to each score.
4. Score each candidate from 1 to 5 on every criterion, using the anchors below. Have two or three people score independently, then discuss any gap of two points or more.
5. Read the ranking and the cut-off flags in the spreadsheet. Then read the section on overrides before you commit.
6. Record the decision, the scores and the evidence gaps. Set a date, usually 60 to 90 days out, to rescore with what you have learned.

The eight criteria

Each criterion is scored from 1 to 5, where 5 is always the most attractive outcome for you. That matters for competition intensity, where a 5 means the field is open or beatable, not crowded. The default weights in the template reflect what most often decides whether a niche works for an early or

expanding product: the pain has to be real, the buyers have to be reachable and someone has to be paying for the problem already.

Criterion	Default weight	Core question
Pain severity	5	How much does this problem cost the people who have it?
Frequency	3	How often does the problem occur?
Reachability	4	Can you find and contact these buyers at a reasonable cost?
Willingness to pay	4	Will they pay a price that supports your business model?
Evidence of spend	4	Are they already spending money or staff time on this?
Unique advantage	3	Do you have an edge here that others lack?
Competition intensity	2	Can you win against the current alternatives?
Expansion path	2	Does winning here open adjacent markets or larger accounts?

Pain severity

A 1 is an annoyance people mention when asked but have learned to live with. A 3 is a problem that costs measurable time or money and appears on someone's list of things to fix this year. A 5 is a problem tied to revenue, compliance, customer loss or a named executive's goals, where people describe workarounds with visible frustration. Gather evidence by asking people to walk you through the last time the problem happened, what it cost and who noticed. Pain that people can only describe in the abstract usually scores lower than they claim.

Frequency

A 1 is a problem that occurs once a year or less, such as an annual audit or a one-off migration. A 3 is monthly or tied to a regular cycle like month-end close. A 5 is daily or weekly, built into how people work. Frequency drives retention and habit, which matters a great deal for subscription products and for AI tools that need repeated use to show value. Evidence comes from asking how many times the problem happened in the last month, not how often it "usually" happens.

Reachability

A 1 means buyers are scattered, hard to identify and not gathered in any channel you can access. A 3 means you can build target lists, but outreach is cold and response rates are uncertain. A 5 means the niche has dense communities, associations, events, publications or partner channels, and you or your team already have a way in. Test this directly: try to book ten conversations in two weeks. How hard that is tells you more than any market map.

Willingness to pay

A 1 means buyers expect the solution to be free or bundled into something they already own. A 3 means they will pay, but at a price that strains your model or needs a long approval chain. A 5 means a buyer with budget authority can say yes at a price that supports healthy unit economics. Evidence

comes from price conversations, not hypotheticals: ask what they would cut to pay for this, who signs, and what they paid for the last tool that solved a similar problem.

Evidence of spend

A 1 means nobody is spending anything on the problem today, which often means it is not a priority. A 3 means there is indirect spend, such as staff time on manual work, spreadsheets or generic tools stretched to fit. A 5 means there is direct spend: an existing vendor, a contractor, an agency or dedicated headcount. Existing spend is the strongest single signal that a budget can be redirected to you. Look for job postings, vendor contracts mentioned in calls, and line items people can name.

Unique advantage

A 1 means you bring nothing a competent new entrant would not also have. A 3 means you have one meaningful edge, such as domain knowledge, a partnership or relevant data. A 5 means you have several compounding advantages: credibility in the community, proprietary data, distribution, or a technical capability that is hard to copy. For AI products, be honest here. Access to a capable model is not an advantage, because everyone has it. Workflow knowledge, data rights and trust usually are.

Competition intensity

A 1 means the niche is crowded with well-funded vendors that customers are broadly satisfied with. A 3 means there are competitors but clear gaps, such as poor fit for a sub-segment, high prices or weak service. A 5 means the current alternatives are manual work, generic tools or vendors customers openly complain about. Remember that "no competitors" can also mean "no market", which is why this criterion carries a lower weight and why evidence of spend sits alongside it.

Expansion path

A 1 means the niche is a dead end: winning it leads nowhere else. A 3 means there are adjacent segments with similar needs that you could reach with moderate product change. A 5 means winning this niche gives you references, data or product depth that carry directly into larger accounts or neighbouring markets. A niche you can dominate is worth more when it is the first step on a path rather than the whole journey.

Gathering evidence

The scorecard is only as good as the evidence behind it. Treat every score as a claim and write its source next to it in the spreadsheet notes. Useful sources include structured customer conversations, sales and support records if you already have customers, job postings that describe the problem, public pricing of existing vendors, community forums and review sites, and your own attempts to reach the buyers.

Keep a simple confidence marker for each score: high when it rests on direct conversations or data, low when it rests on assumption. A candidate with several low-confidence scores is not a weak candidate. It is an unknown one, and the right next step may be a two-week research sprint rather than a decision.

Worked example (hypothetical)

The following is a hypothetical comparison for an imagined company building an AI assistant that reads and processes incoming documents. The team is choosing between three niches. The names and scores are illustrative, not drawn from real companies or research.

Criterion (weight)	A: Mid-size accounting firms, client document intake	B: Independent dental clinics, AI call handling	C: In-house legal at growth-stage SaaS, contract review
Pain severity (5)	4	4	4
Frequency (3)	5	5	3
Reachability (4)	4	2	3
Willingness to pay (4)	4	3	5
Evidence of spend (4)	4	4	5
Unique advantage (3)	3	2	2
Competition intensity (2)	3	2	1
Expansion path (2)	4	3	4
Weighted total (max 135)	**106**	**87**	**97**
Percent of maximum	**78.5%**	**64.4%**	**71.9%**

With a cut-off of 65%, the scorecard ranks A first, C second and flags B as below the cut-off.
The story behind the numbers matters more than the totals. Accounting firms process client documents every day, the team has a former practice manager who can open doors through professional associations, and firms already pay for intake software and junior staff time. Dental clinics have real and frequent pain, but the team has no route into the market and several call-handling vendors already sell there. In-house legal teams pay well and already spend on outside counsel and contract tools, but the competitive field is crowded and the team brings no particular credibility to lawyers.

How to read the result

Start with the gap between candidates, not the ranking. A difference of five percentage points or less is within the noise of honest scoring; treat those candidates as tied and decide on evidence quality or strategic fit. A gap of ten points or more usually reflects a real difference.

Then look at the shape of each score. Two candidates with the same total can be very different bets. One that scores 5 on pain and 2 on reachability is a distribution problem; one that scores 3 across the board is a niche without a strong reason to win. The Excel scorecard shows weighted contributions per criterion so you can see which factors carry each total.

Finally, look at the cut-off flags. A candidate below the cut-off should not be chosen without a written reason. The flag does not forbid the choice; it requires you to explain it.

When to override the scorecard

The scorecard supports judgement; it does not replace it. There are legitimate reasons to choose a lower-ranked candidate.

- **A single criterion is disqualifying.** If a niche scores 1 on willingness to pay, a high total elsewhere does not rescue it. Some teams set hard floors, such as no score below 2 on pain or spend.
- **Your evidence is lopsided.** If the leader rests on assumptions and the second rests on twenty conversations, the second may be the safer bet.
- **Strategic constraints apply.** A regulatory barrier, a partner commitment or a funding milestone can legitimately change the answer.
- **The team cannot sustain the work.** A niche the team finds tedious will be served poorly. Interest is not a criterion, but its absence is a risk.

When you override, write down why, and name the evidence that would prove you wrong. That keeps the override honest and gives you a clear trigger to revisit it.

Common mistakes

- **Defining niches too broadly.** "SMBs" or "enterprise" cannot be scored meaningfully because the answers vary too much inside them.
- **Scoring the market you wish existed.** Scores should reflect what buyers say and do now, not what they might do once they see your product.
- **Changing weights after scoring.** Adjusting weights to make a favourite win defeats the purpose. Fix weights first.
- **Treating no competition as good news.** An empty field often means no budget. Check evidence of spend before you celebrate.
- **Counting model access as an advantage.** For AI products, the model is a commodity. Your edge lies in data, workflow fit, trust and distribution.
- **Scoring alone.** One person's scores mostly reveal that person's preferences. Independent scoring by two or three people surfaces disagreement worth discussing.
- **Never rescoring.** A niche decision made on week-one evidence should be checked against what you learn in the following months.

When not to use this

Do not use the scorecard when you have a single candidate and the real question is whether to pursue it at all. That calls for a go or no-go test with explicit success criteria, not a comparison. It is also the wrong tool once you already have strong traction in one segment; at that point your customer data is better evidence than any scoring exercise. And if you cannot yet name the specific buyers in each candidate niche, put the scorecard aside and do the research first. Scoring vague candidates produces precise-looking numbers with nothing underneath.

Editable template: Scorecard (Excel), available at ratomir.com/resources/niche-evaluation-scorecard

© 2026 Ratomir Jovanovic. Free to use inside your team. Please link to ratomir.com/resources/niche-evaluation-scorecard when you share it.