

HEARD OF
BY CONTENT ALLIES

Audience before outreach.

A practical B2B playbook for paid podcast distribution,
audience building, and account-based marketing.



A 90-DAY PILOT. CLEAR MEASURES. USEFUL NEXT STEPS.

2026 EDITION / FOR MARKETING EXECUTIVES

Give the podcast a commercial job.

Paid podcast distribution can be the beginning of an audience strategy. Its value depends on the market you choose, the content you distribute, and the evidence you collect after delivery.

This guide proposes a practical pilot for B2B marketing teams: choose one buying problem, promote a useful episode, extend its best ideas into other media, and create a voluntary path to a deeper interaction. Treat the first campaign as a test of a defined audience and message, not a shortcut to pipeline.

Three decisions to make before spending

- 1. Who should care?** Name a segment and the people involved in a buying decision. An account list can focus the program, but every channel will have different targeting capabilities.
- 2. What should they learn?** Choose a subject your experts can explain with enough depth to be useful. Connect it to a real business problem rather than forcing every episode to be a product pitch.
- 3. What would justify another flight?** Set delivery, quality, and learning criteria in advance. A lower cost per download is not sufficient if the audience or content is a poor fit.

The measurement boundary

A download is evidence of file delivery under a counting method. It does not establish that a named person or target account listened, that the episode was completed, or that a sale resulted. Keep those claims separate. [1, 2, 6]

Use this guide with: marketing leadership, content owners, paid media, marketing operations, and a sales counterpart. The worksheets are proposed operating tools, not industry benchmarks or a claim of proven campaign results.

Four jobs. One commercial story.

01 Choose the market / GTM

Set the segment, buying problem, offer, account priorities, and exclusions. Write the question that the pilot should answer.

Working output: A market and account brief.

02 Build the audience / demand generation

Use paid podcast promotion to introduce a useful point of view to relevant audiences. Give the episode a defined media budget and delivery objective.

Working output: An audience-growth campaign.

03 Deepen engagement / media distribution

Turn the episode's strongest ideas into approved clips and practical resources. Give each format and channel a distinct job in the sequence.

Working output: A coordinated media plan.

04 Connect account action / ABM

Use supported company-level engagement and voluntarily identified contacts to inform relevant follow-up. Combine this with sales feedback and CRM evidence.

Working output: An account learning review.

These are related activities, not interchangeable metrics. Podcast audience growth can support an account strategy without providing account-level proof of listening.

Start with fit. Then check reachability.

Use evidence from customers, lost deals, and the sales team to choose a cohort with a shared problem. Demandbase’s account-selection guidance similarly starts with customer characteristics and the buying committee. Its platform-specific recommendations are not prerequisites for this pilot. [5]

Decision	Write down	Why it matters
Market fit	Sector, company size, geography, and use case.	Keeps the creative and audience brief coherent.
Commercial fit	The problem you solve and the offer this supports.	Connects audience building to the GTM plan.
Buying group	Champion, evaluator, economic buyer, and blockers.	Reveals which perspectives the content needs.
Account cohort	Named companies, rationale, exclusions, and owner.	Makes marketing and sales priorities explicit.
Reachability	Channel capability, match coverage, and audience constraints.	Prevents account-level promises on unsuitable inventory.

Keep the pilot narrow enough to learn

Choose a cohort your team can review account by account. The right size depends on market concentration, available inventory, platform eligibility, and sales capacity; do not force an arbitrary account count. A small named list may require broader professional targeting for podcast discovery.

Review exclusions before activation: competitors, irrelevant territories, incompatible use cases, and accounts already handled through a different commercial program. Keep a dated copy of the final cohort so later reporting uses the same denominator.

Give every role a useful next step.

Begin with one question a buyer is already trying to answer. Record an episode that explains the tradeoffs. Then identify the parts worth distributing separately. The table is an illustrative manufacturing example, not a customer campaign.

Role	Question worth answering	Content and next step
Operations leader	Where do changeovers create avoidable disruption?	Episode on diagnosing bottlenecks; optional workflow checklist.
Technical evaluator	Will the approach fit our operating environment?	Short technical clip; approved application note.
Finance / economic buyer	What assumptions should we use to evaluate the case?	Brief on evaluation inputs; conversation about the buyer's own context.
Procurement	What should we verify about a supplier?	Supplier-evaluation checklist; explicit request for a supplier briefing.

Brief the asset before editing it

For each clip or resource, record the intended role, question, claim, supporting evidence, destination, and reviewer. Preserve the context behind any numbers or performance statements. A short excerpt should not make a stronger claim than the full conversation.

A useful sequence: introduce a problem, explain a decision, offer a practical tool, then invite a relevant conversation. The buyer chooses whether to take each step.

Match the media to the job.

Use a shared audience brief across the program. Do not assume the same audience can be identified or reached in every environment. LinkedIn offers company-list matching and professional targeting; that capability does not transfer automatically to podcast placements. [3]

Media	Job in the pilot	What to confirm
Paid podcast promotion	Introduce your own show and acquire downloads.	Placement type, listener journey, targeting, download definition, quality controls.
Professional social	Distribute a relevant clip or resource; use account matching where eligible.	Match coverage, role filters, minimum audience rules, exclusions, reporting availability.
Publisher / contextual media	Reach people around a relevant business subject.	Actual inventory, context, format, geography, frequency, and data available.
Owned resource page	Offer a checklist, guide, or requested briefing.	Clear CTA, functioning analytics where permitted, form handling, and next-step owner.

Ask these questions of every distribution provider

Where will the content appear? What action starts playback or a download? How are invalid and duplicate events handled? Which targeting attributes are observed versus inferred? What can the report show? What happens if delivery quality falls outside the agreed terms?

This program buys attention for your own content. It is different from buying a host-read sponsorship for your company on another publisher's podcast, and from paying for hosting or RSS directory submission.

Three layers. Three kinds of evidence.

Layer	Useful measures	Boundary
Delivery	Downloads, placements, spend, cost per download under agreed rules.	Not proof of completion, buyer identity, or incrementality.
Engagement	Platform consumption, landing-page actions, resource requests.	Different sources cover different audiences; no automatic cross-channel deduplication.
Account / business	Supported company engagement, known contacts, accepted conversations, CRM progression.	An observed association does not establish that the media caused the outcome.

IAB Tech Lab describes podcast measurement based on server logs. Apple Podcasts separately provides aggregated consumption analytics for activity inside its platform. LinkedIn’s company reporting covers engagement with content in its own environment. These are different observation systems. [1, 2, 4]

Write a measurement dictionary

For every KPI, specify the event, source system, filtering rule, time window, denominator, owner, and limitation. Record whether the measure is campaign-attributed, all-show, account-matched, self-reported, or modeled. Do not relabel all-show downloads as paid campaign delivery.

Example reporting language

“The campaign delivered X measured downloads under the agreed definition. Separately, Y known contacts requested the guide. We cannot establish that those contacts are the people represented by the downloads.”

Ask the measurement provider which guideline version and certification scope apply. IAB’s standards page reviewed for this guide linked v2.2; its July 2026 announcement described v2.3 as a public-comment release. Do not imply certification or finalized adoption from that announcement alone. [1, 6]

Buy a useful test, not a big total.

Agree separate budgets for media, content adaptation, setup, and reporting. Write down whether a quoted cost per download includes service fees. The calculation only has meaning when its numerator and denominator are defined consistently.

Illustrative arithmetic, not a price or forecast

If a campaign spends \$6,000 on media and reports 4,000 attributable valid downloads, media CPD is \$1.50. If scoped service fees add \$2,000, fully loaded cost is \$2.00 per download. Neither number is a cost per buyer or per lead.

Design the comparison before the campaign

Record a baseline for the same show, account cohort, and definitions. If feasible, hold out a comparable account group from account-targeted media and compare changes over the same period. Other exposure can reach either group; small cohorts, sales activity, and seasonality can distort the result. Treat weak comparisons as directional learning.

Decide what changes, and why

Scale: delivery is credible, the audience and engagement evidence fit the objective, and the team can support the next step. Increase investment in a defined increment, then review again.

Revise: media delivers but the subject, destination, or engagement is weak. Change one major variable where practical, documenting what changed and when.

Stop: measurement is unreliable, placements do not fit, spend exceeds the agreed ceiling, or the campaign cannot answer the learning question. Pausing is an acceptable outcome of a pilot.

Set numeric thresholds with your team before launch using your own economics and baseline. This guide intentionally supplies no universal CPD, conversion, or pipeline benchmark.

A working cadence for the team.

When	Work and owner	Exit condition
Days 1-15	Marketing lead + sales: choose the market, problem, cohort, baseline, and handoff.	One approved brief and a written learning question.
Days 16-30	Content + media + operations: approve assets, inventory, tracking, budget, and reporting definitions.	Destinations tested; owners and launch criteria agreed.
Days 31-60	Media lead: launch controlled flights; inspect delivery and quality. Content owner: review engagement.	Weekly decision log with issues, changes, and rationale.
Days 61-90	Marketing + sales: assess audience evidence, account activity, known inquiries, and economics.	Documented scale, revise, or stop decision; next test defined.

Make the weekly review short and consequential

Use five questions: Are we delivering to the planned environment? Are the measurements behaving as expected? What content earns further engagement? What known account activity is relevant? What single decision should change next week’s plan?

Define the handoff before the first inquiry

A requested briefing can go to an assigned owner with the topic, stated need, and permission context. Passive exposure should remain a media signal. Do not tell a prospect “we saw you listen” when the evidence is only anonymous delivery or company-level activity.

The 90-day cadence is a proposed operating schedule. Some business outcomes will take longer to observe; a pilot can end with useful audience learning before it produces attributable revenue.

Keep the framework. Change the context.

Use these examples to adapt the pilot. They describe possible programs, not sector-wide benchmarks or observed customer results.

Private equity and portfolio companies

Pilot with one portfolio company and one commercial market. Standardize briefs and reporting across the portfolio only after the first team can use them. Keep company budgets, audiences, and outcomes separate; the program is commercial marketing, not investor solicitation.

Consulting and professional services

Organize around a practice-specific decision. An expert episode can lead to a practical diagnostic or briefing. Give business development enough topic context to offer a relevant conversation, while protecting confidential client material.

B2B healthcare

Choose professional audiences and a business problem such as workflow or implementation. Assign reviewers to product and clinical claims before activation. Do not infer a person's health status from engagement; this is not a patient-acquisition program.

Manufacturing and industrial B2B

Start with an application and its buying group. Pair the episode with approved technical material. Separate operations, engineering, and purchasing questions, and keep performance claims tied to the conditions in which they are valid.

Common deliverables: an audience brief, approved content, channel plan, measurement dictionary, review cadence, and a clear owner for the next voluntary interaction.

Write the pilot on one page.

Use this in the kickoff. Fill the gaps before approving media spend.

Prompt	Your team's answer
Market and commercial offer	Segment: _____ Offer / use case: _____
Cohort and buying group	Account list owner: _____ Priority roles / exclusions: _____
Learning question	We want to learn whether _____ _____
Content and destination	Episode / core idea: _____ Companion asset / next step: _____
Media and economics	Channels / targeting limits: _____ Media budget / fees / ceiling: _____
Measurement	Sources / definitions / baseline: _____ Reporting window / owner: _____
Decision rules	Scale when: _____ Revise or stop when: _____
Handoff and review	Inquiry owner: _____ Sales context / review date: _____

Add the approved account list and a channel-level plan as attachments. This worksheet is intentionally static and printable; it does not collect or send your answers.

Make the decision visible.

Review field	Record before the next flight
Delivery and economics	Spend; valid campaign downloads; media and fully loaded CPD; quality exceptions.
Audience engagement	Source-specific consumption and actions; coverage limits; content worth testing again.
Known account activity	Matched platform evidence, voluntary inquiries, accepted conversations; method and window.
Decision and owner	Scale / revise / stop; rationale; next hypothesis; budget; accountable person and review date.

Method note. This guide combines public primary-source documentation with an original proposed planning framework. Examples, arithmetic, timelines, and worksheets are illustrative. It contains no original campaign-performance study and makes no outcome guarantee. Sources checked September 24, 2026.

[1] [IAB Tech Lab - Podcast Measurement Technical Guidelines](#)

[2] [Apple Podcasts for Creators - Analytics](#)

[3] [LinkedIn - Company Targeting](#)

[4] [LinkedIn - Measure engagement from your key companies](#)

[5] [Demandbase - How to identify accounts for ABM \(Jan. 8, 2026\)](#)

[6] [IAB Tech Lab - v2.3 public-comment announcement \(July 21, 2026\)](#)

Heard Of by Content Allies

Start with paid podcast distribution. Build a wider media strategy around the audience you want to earn. Bring this worksheet to your next planning conversation.