LEAD GENERATION AND APPOINTMENT SETTING

From search evidence to signed retainers

Lead Generation and Appointment Setting for SEO Agencies

Find businesses with a search-dependent buying journey, an addressable decision and a route to implementation. Use evidence to earn a useful first review.

26	12	06
numbered sections	**editable emails**	**public workflows**

Research edition: 30 October 2026. Eight working templates, model calculations, source notes and a separate publishing pack accompany the guide.

Published research and public workflow observations are attributed. Commercial, production and test examples are fictional. No named company is presented as a client or a qualified prospect.

Inside the guide

Choose the section that matches the decision you need to make. The full guide connects account research, first offers, appointments, delivery and economics.

01. Define the pipeline before choosing a tool

02. What the research establishes, and what it does not

03. Select an account cell with implementation potential

04. Use six public workflows to improve the questions

05. Separate a fact, a diagnosis and a commercial claim

06. Build the list around search economics and ownership

07. Choose a trigger that changes the decision

08. Build first offers that lead naturally to paid work

09. Run a public sample that another analyst can reproduce

10. Diagnose falling traffic without an AI shortcut

11. Handle query coverage and branded splits correctly

12. Connect search to buyer economics without inventing revenue

13. Scope the implementation path during discovery

14. Write the campaign around one commercial question

15. Use a sequence that adds information

16. Turn interest into an agreed search review

17. Accept, hold or decline the first meeting

18. Use proof that matches the buying situation

19. Price acquisition against contribution and capacity

20. Test one uncertainty at a time

21. Troubleshoot the pipeline at the failing stage

22. Run the first 90 days with decision gates

23. SEO campaign message bank

24. Working templates

25. Sending and measurement operating notes

26. A campaign is ready when these decisions are concrete

Source notes follow the numbered sections. The editable Markdown, structured JSON, message bank and research registers are included in the content package.

01. Define the pipeline before choosing a tool

Use these definitions consistently. A prospect is a distinct buying entity that passes the campaign's eligibility rules. A contact is a person at that entity. A reply is a human response. A positive account is an eligible account where someone expresses relevant interest. An appointment is an agreed conversation with a date, time and purpose. An accepted held meeting is a conversation that occurred and meets the written first-discussion criteria. An opportunity is a separately qualified commercial project.

Measure	Operational definition	Common counting error
Eligible accounts reached	Distinct eligible entities with at least one valid contact attempt	Counting every contact as another company
Positive-account rate	Positive accounts / eligible accounts reached	Counting referrals and repeat replies twice
Booking rate	Agreed appointments / positive accounts	Treating a resource request as calendar consent
Held rate	Meetings that occurred / appointments due	Counting future meetings as no-shows
Acceptance rate	Accepted held meetings / held meetings reviewed	Counting anyone who attends
Opportunity rate	Qualified projects / accepted held meetings	Calling every discovery conversation pipeline
Win rate	Won projects / decided proposals in the named cohort	Including open proposals as losses

Fix the unit of analysis before comparing campaigns. Keep an account ID, contact ID, conversation ID and meeting ID. A second contact can help establish ownership without creating a second account-level success.

Decision: optimise for suitable commercial conversations and contribution, with sending activity as an operating input. Owner: the agency leader and sales owner approve the definitions together before outreach starts.

02. What the research establishes, and what it does not

The research points toward more careful diagnosis, not a universal script about SEO disappearing.

Evidence	Population and period	Useful interpretation	Limit
Pew: 8% traditional-result clicks with an AI summary versus 15% without	900 US adults; March 2025 browsing; results reconstructed in April	Search context can change the chance of a result click	Observational, not a site-level experiment
Ahrefs: roughly 58% lower top-position desktop CTR relative to its modelled counterfactual	300,000 keywords; December 2023 and 2025	Rankings alone do not explain click performance	Not a blanket 58% traffic loss
SparkToro: 68.01% zero-click share	US browser-based mobile and desktop sample, January-April 2026	A website session captures only part of search behaviour	Different historic panels; zero click is not necessarily zero business value

The Ahrefs estimate is not simply the raw change from a 7.3% CTR to 1.6%. Its method also adjusts for the decline in the comparison group. Pew measured different searches and behaviours. Do not average the two studies or turn either into a prospect's lost-revenue estimate. S01 S02 S03

Google says existing search foundations remain relevant to its AI features, with no special technical requirement or AI-only markup. This guidance concerns Google AI Overviews and AI Mode, not every AI product. Eligibility still does not guarantee display. S04

The practical offer is a bounded diagnosis: which queries, pages, markets and conversion paths explain the buyer's situation? The next step is not automatically a site-wide content programme or a new label for ordinary SEO work.

03. Select an account cell with implementation potential

An account cell combines one business model, one search job, one decision owner and one type of first offer. Geography and delivery capability narrow the cell further. Industry alone is too broad.

Account cell	Search job	Owner to verify	First offer	Poor-fit condition
Multi-location services	Convert service-area discovery into serviceable enquiries	Central marketing and local operators	Five-location ownership sample	Nobody can change local pages or follow up leads
Ecommerce	Help shoppers discover available, profitable products	Ecommerce and merchandising	Category/filter sample	Out-of-stock or margin issues dominate the opportunity
B2B software	Connect use-case searches to accepted opportunities	Growth and product marketing	Query-to-opportunity map	No agreed qualified outcome or implementation owner
Professional services	Reach buyers of a defined, valuable service	Practice lead or marketing	Service-page and enquiry review	Unclear service area, offer or capacity
Migration-stage business	Preserve and improve a changing site	Digital owner and engineering	URL and launch risk map	No access to release planning
Publisher or education platform	Match useful content to audience and commercial purpose	Editorial and audience leads	Content-cohort review	Traffic is the only monetisation assumption

Start with a cell where the agency has relevant delivery skill and proof. A technical SEO team should not sell a large digital PR programme because it appears easier to pitch. A content agency should not imply it can own a complex platform migration without a qualified delivery partner.

Make the buyer's constraints part of eligibility. Language coverage, content review, development access and regulated subject matter can change the work more than company headcount.

04. Use six public workflows to improve the questions

These are purposive examples of page and business models. They are not clients, performance audits or a list of companies seeking an agency. No private analytics were accessed.

Public example	Observed workflow	Commercial question	Appropriate first asset
Wise currency converter	Utility content with a money-transfer action	Which utility-page cohort connects to a valuable customer action?	Query-to-conversion map
Zapier app directory	Structured discovery of integrations and products	Who can change a whole page family?	Template and owner map
Ruggable all-rugs category	Product navigation by room, size and collection	Which category combinations are useful and commercially viable?	Category/filter sample
The Grounds Guys location directory	Local service discovery across locations	Who owns the page, profile and serviceable enquiry?	Location ownership map
Trainline journey pages	Route discovery and ticket-booking path	How do route demand and live availability affect bookings?	Route-page cohort brief
Atlassian migration content	Education for a product migration decision	Which searches lead to an accepted migration discussion?	Content-to-opportunity map

The source register preserves the URL, observation, research date, unanswered question and asset suggestion for every row. Search rankings, traffic, agency relationships and budgets remain unknown. A visible page model tells you what to investigate. It does not prove weak SEO.

For a real campaign, use comparable smaller or larger accounts only where the agency can serve their operating model. A well-known example makes a workflow easy to understand; it does not establish that the brand is a sensible cold-outreach target.

05. Separate a fact, a diagnosis and a commercial claim

Store these in different fields. A fact is observable: a company published a migration announcement, or a category page exposes filter combinations. A diagnosis is an interpretation supported by analysis: a sampled set has inconsistent canonical signals. A commercial claim estimates an outcome: resolving those signals may support a defined search or conversion goal.

Evidence level	Example	What can appear in outreach
Public fact	A dated announcement names a new market	The fact, date and a relevant question
Reproducible sample	Five URLs show the same observable issue	Exact scope and how it was checked
Buyer-confirmed condition	The buyer reports a release or measurement problem	Accurate summary with permission to use it
Commercial hypothesis	Better coverage may create more qualified demand	A labelled hypothesis to test

Never say "you are losing USD 40,000 a month" because a third-party tool estimates traffic. Never diagnose a penalty because traffic fell. Never claim a search opportunity is "low competition" without defining the country, intent, query set and date.

A usable evidence row contains the source URL, capture date, market/device context where relevant, exact observation, sample boundary, inference, unknowns and next verification step. Stop research when the first offer is specific enough to be useful. More screenshots do not automatically create a stronger case.

06. Build the list around search economics and ownership

The minimum research record should answer: what does this company sell, where does search help the buying journey, who owns that journey, what can the agency deliver, and what decision would justify a conversation?

Use headcount or estimated traffic as a routing aid, not a release rule. A small specialist with valuable transactions can be a better fit than a large company whose buying model depends on partnerships. A high-traffic publisher may have little ability to fund the agency's proposed work.

Release condition	Confirm now	Can remain unknown until discussion
Business fit	Relevant offer and market	Full revenue breakdown
Search fit	A defined page family or search-dependent journey	Exact first-party traffic
Service fit	Agency can deliver the proposed first asset	Complete implementation backlog
Role fit	Contact plausibly owns or routes the decision	Entire buying committee
Contact route	Appropriate route and no stop instruction	Future procurement timetable

For one-to-one research, inspect the company's own website, relevant public announcements and professional role information. For a segmented campaign, use a verified common page model and offer. Do not insert "I audited your site" into automation unless someone completed the described review for that account.

Review a suggested pilot batch of 25 account records manually. This is an operating choice, not a response benchmark. Record why each account passes, is held or is excluded. Fix the rule when reviewers disagree repeatedly.

07. Choose a trigger that changes the decision

A timely fact is useful only when it changes the proposed work. A migration, new market, new service line, acquisition or known contract-review date can change scope and ownership. A job post may help locate a team but does not establish agency budget or a request for proposals.

For a migration signal, ask whether the release plan already includes URL mapping and search monitoring. For a new location, ask how the company assigns ownership of pages, profiles and enquiries. For a new product

category, ask which commercial pages should represent the range and who can implement them.

Trigger	Relevant first output	Unsupported inference to avoid
Confirmed platform migration	Launch risk and dependency map	The current site is failing
New geography	Locale, language and conversion-path brief	The company needs international SEO now
New locations	Page/profile/lead ownership map	Every location has ranking problems
New category	Sample taxonomy and availability map	Search demand guarantees profitable sales
Growth-role hire	A role-specific process question	Hiring means external vendor intent

Date the trigger and preserve the exact source. If it is old, call it context. Ask whether the review is still relevant instead of manufacturing a deadline.

08. Build first offers that lead naturally to paid work

A good first offer has a decision, a bounded input, a tangible output and an honest limit. It demonstrates how the agency thinks without giving away an undefined consulting project.

Offer	Input	Output	Boundary	Natural paid next step
Category opportunity sample	One category and five public URLs	Useful page/filter questions and implementation gaps	No revenue forecast	Full taxonomy and implementation scope
Search-to-pipeline map	One service and agreed conversion stages	Query/page/stage measurement dictionary	No private account required to start	Measurement setup and focused SEO plan
Migration readiness brief	Announced release and sample old/new URLs	Dependencies, checks and owners	Not launch sign-off	Migration support with engineering
Local enquiry ownership sheet	Five agreed locations	Page/profile/lead responsibility map	No guaranteed map position	Location rollout and measurement
Content-cohort review	A defined topic or page family	Keep, update, consolidate or investigate decisions	No automatic deletion list	Editorial and implementation programme

State the input you need and the output you will send. If the prospect requests the asset, deliver it before asking for a meeting. A brief can contain one specific unresolved decision that makes a discussion useful. It should also work as a standalone document.

Do not describe a publicly visible error as the only important SEO issue. It may be low impact, deliberate or already scheduled for repair. The commercial discussion begins by establishing consequence and ownership.

09. Run a public sample that another analyst can reproduce

Choose a small, relevant set of URLs rather than claiming a complete audit. Suggested starting scope: one page family, five examples, one market and one commercial action. Log the selection rule so the sample is not just the five worst pages someone could find.

Record the page purpose, visible offer, internal discovery route, destination action, observed technical questions and access limitations. If a crawler, analytics account or search tool was not used, say so. A blocked request is missing evidence; it is not proof that Google cannot crawl the site.

Use a short evidence card:

- Observed: the specific URL and condition, with the date and context.
- Scope: the page family and the number of examples reviewed.
- Why it may matter: the customer task or implementation dependency.
- Unknown: current performance, intentional design, backlog and owner.
- Next check: the smallest piece of first-party evidence needed.

The fix should be proportionate to confidence. "Verify the canonical and indexing policy for this filter family" is better than "delete 10,000 pages" when the agency has not reviewed demand, inventory or indexation.

Google's documentation provides the relevant technical foundations and policy boundaries. A checklist does not replace interpreting the buyer's actual site. S05 S10

10. Diagnose falling traffic without an AI shortcut

Start with measurement integrity and comparable periods. Confirm property, timezone, search type, country, device, branded classification, page cohort and settled data. Compare equal-length periods and, where relevant, year-over-year seasonality. Preserve known releases and tracking changes.

Then separate demand, visibility, click behaviour and conversion. Falling impressions can reflect demand, coverage or eligibility. Stable average position does not prove stable ranking for the same queries. Falling CTR can reflect query mix, features, snippets, competitors or device mix. Falling enquiries can come from a form or offer change after the click.

Pattern	Investigate before making a claim
Impressions down, position apparently flat	Query mix, demand, coverage, aggregation
Impressions stable, CTR down	SERP context, page/query mix, titles and snippets
Clicks stable, accepted leads down	Tracking, forms, offer, qualification and sales follow-up
Brand clicks up, non-brand down	Brand definition, campaign spillover and page cohort
One template falls, rest stable	Releases, rendering, internal links and indexation

The next deliverable is a diagnosis tree with evidence and owners, not an automatic content order. Google defines CTR as clicks divided by impressions and average position through its own aggregation rules. Respect those definitions when building the comparison. S06

11. Handle query coverage and branded splits correctly

Search Console query rows are not the same thing as a complete record of every search. The API does not guarantee all rows. Page and property groupings also differ. Keep the property total, exported row total and known coverage limitations visible. Do not force them to reconcile by inventing an "other" segment with a false query identity. S06 S07

Worked example, fictional: property clicks total 10,000; the chosen query export contains 8,200. Of those rows, 2,200 clicks match the agreed brand rule and 6,000 do not. The correct claims are "2,200 observed brand clicks" and "6,000 observed non-brand clicks, with 1,800 property clicks not allocated through this export." Observed non-brand share is 6,000 / 8,200 = 73.17% of represented query clicks. It is not automatically 60% of a completely classified property.

Define brand variants before measuring: company names, domain variants, product brands, misspellings and ambiguous words. Review a sample of matches and non-matches. A dictionary word in a brand name can make a broad regular expression overcount.

Save the classification version, date, search type and aggregation. When the definition changes, restate the earlier period or disclose that the comparison changed. This protects both the agency's analysis and the promises made during appointment setting.

12. Connect search to buyer economics without inventing revenue

Build the commercial model from the buyer's own funnel. For a service business, use sessions from the agreed page cohort, enquiry rate, qualified-enquiry rate, won-client rate and contribution per new client. For ecommerce, use net orders, returns, fulfilment costs and contribution. For B2B, use accepted opportunities and sufficiently mature sales cohorts.

Fictional service example:

Driver	Assumption	Calculation
Additional monthly visits	2,000	Scenario, not a forecast
Enquiry rate	2%	40 additional enquiries
Qualified share	50%	20 qualified enquiries
Qualified-to-won rate	20%	4 expected clients
Contribution per client	USD 1,000	USD 4,000 contribution
SEO and implementation cost	USD 3,500	USD 500 net scenario contribution

If the additional visits are only 1,000, the same funnel gives USD 2,000 contribution and a USD 1,500 loss after cost. If the qualified share falls, the result also changes. Show that downside during scoping.

This model does not establish that SEO caused the extra visits, that the visit increase will occur, or that delivery produces immediate steady-state returns. Put ramp, implementation delay, sales cycle and attribution limits beside the result. Value the decision honestly before selling an expensive programme.

13. Scope the implementation path during discovery

Ask who can change templates, navigation, redirects, page copy and conversion events. Establish release cadence, review requirements, access and rollback procedures. A recommendation that cannot be implemented is not a credible growth plan.

For a migration, map old and new URLs, relevant signals, assets, analytics and ownership. Google warns that significant moves can involve ranking fluctuations while pages are recrawled and reindexed. A readiness brief should expose dependencies without promising a fixed recovery date. S08

For local work, establish which entity controls the Business Profile, the landing page and lead follow-up. Search results depend partly on the searcher's location, so a screenshot from the agency's office cannot demonstrate the whole service area. S09

For content, identify the subject expert, legal or product reviewer, publication owner and update trigger. For technical work, identify the engineering counterpart and an accepted test. For digital PR, define editorial fit, research value and reporting rather than a guaranteed number of ranking gains.

End discovery with an implementation map: requirement, owner, dependency, effort range, evidence, review date and decision. Unknown effort belongs in a scoped assessment, not a fixed retainer sold on hope.

14. Write the campaign around one commercial question

Use the LS Cold Email structure: why this buyer, why this fact matters now, what useful output the agency can provide, and one clear next step. A current trigger is optional when none is verified. Relevance can stand on a real workflow.

Manual route: use one or two verified details from the company and role. Segmented route: use a truthful shared page model and a relevant offer. Do not generate synthetic compliments or private-performance claims.

Keep the first message below 100 words including the subject, signature and required footer. Prefer around 80 or fewer before replacing placeholders. Use an honest subject. The last request should be easy to answer and immediately precede the signature. Do not combine a calendar request, a PDF request and a referral request in one email.

The campaign brief should contain the account cell, evidence rule, contact role, first asset, supported capabilities, exclusions, response owner and acceptance criteria. The list and offer create much of the relevance before wording begins.

A campaign is ready when another researcher can apply its rules and another salesperson can explain the promised artifact. If only the founder understands the offer, simplify it before increasing volume.

15. Use a sequence that adds information

The following cadence is an illustrative operating plan. It is not a proven universal optimum. Follow the recipient's response and applicable rules; stop when asked.

Touch	Example timing	New value	Decision
Initial	Day 1	One verified fact and one asset	Does the review concern them?
Follow-up	Day 4	One precise scope or useful example	Is the asset worth receiving?
Follow-up	Day 9	Explain the evidence or implementation limit	Is there a suitable owner?
Close the loop	Day 16	Clear pause with an option to resume	Continue, defer or stop

Do not repeat "just checking in" while adding no useful information. Do not claim the buyer saw an email because an open event fired. Do not call a resource request agreement to a meeting.

If the buyer replies, end the automated sequence for that conversation and let a named owner respond. Keep cross-contact activity coordinated at account level. One person's deferral should not trigger a different salesperson to restart the pitch the next morning.

Review reply meaning weekly. A higher reply rate driven by objections or wrong-person responses is a signal to repair targeting, not a successful campaign to scale.

16. Turn interest into an agreed search review

Suggested operating target: a human reviews substantive replies within one business hour during stated coverage hours. This is a service target to test against capacity, not a claim that a particular response time causes a given uplift. Auto-acknowledgement is not a substantive answer.

First respond to the actual request. Send the promised asset. Answer the scope question. Then propose one next decision. A useful invitation says what will be reviewed, what input is optional, who should attend and what the buyer will leave with.

Example agenda for a 20-minute first review:

- Minutes 0-3: confirm the service, market and business decision.
- Minutes 3-8: inspect the bounded evidence and correct assumptions.
- Minutes 8-13: identify the owner, implementation path and measurement gap.
- Minutes 13-18: agree whether a focused assessment or programme is justified.
- Minutes 18-20: record the next output, owner and date.

Offer two real times with a timezone after the buyer wants a discussion. A calendar link can be an option. Send an invitation only after agreement. Confirm the purpose and attendees in the invite. Do not treat missing private analytics as a failed first meeting when the agreed first output uses public evidence.

17. Accept, hold or decline the first meeting

An accepted held SEO discussion has a relevant buying entity, an identifiable search-dependent workflow, an appropriate process contact, a real question, an agreed conversation and a feasible agency contribution. Budget, complete technical access and final procurement authority can be established later.

Condition	Accept when	Hold when	Decline when
Workflow fit	Named page family or commercial journey	Buyer says only "more traffic"	Work is outside capability
Ownership	Contact owns or can route the decision	Relevant team unknown	No relevant buying entity
First output	Both sides know the review's purpose	Asset and agenda still vague	Buyer expects a guaranteed ranking
Timing	Current question or agreed future review	No active decision yet	Stop instruction or no permitted route

Keep meeting acceptance and opportunity qualification separate. A valuable exploratory meeting can lead to a paid assessment, a deferred review or a clear no-fit decision. Calling all three "SQLs" hides the actual sales work.

The sales owner records the decision and a short reason within a suggested two business days of the meeting. Use the same rule across team members. Review disputed rejections with evidence rather than retroactively changing the definition to match a sales outcome.

18. Use proof that matches the buying situation

An approved case study should name the client or explain legitimate anonymity, scope, period, starting point, implementation work, measures and limitations. Show contribution from the agency without assigning it every business improvement that occurred.

For a technical case, show the affected template and implemented change. For content, show the commercial purpose and production process. For local work, show service-area and enquiry definitions. For a migration, show launch preparation and monitoring. A logo wall cannot replace those details.

If relevant client proof does not exist, use a reproducible public teardown, a labelled synthetic example or a bounded paid assessment. Public company workflows in this guide demonstrate research method. They do not imply past work with those companies.

Treat competitor comparisons similarly. A third-party traffic estimate can generate a question. It cannot establish competitor revenue, margin or the prospect's exact lost demand. Avoid using another site's estimated traffic as the result the agency is promising to reproduce.

Decision: select the proof after choosing the account cell. Owner: delivery confirms that the promised capability and case interpretation are accurate before sales uses them.

19. Price acquisition against contribution and capacity

Agency revenue is not contribution. Separate the retainer from direct delivery labour, contractors, tools and other costs that increase with that client. Overhead and tax still matter after contribution. Promethean reports 13% average net margin for broad digital agencies in 2025; that is neither SEO-only data nor a contribution-margin assumption. S11

Fictional agency example: 400 eligible accounts yield 20 positive accounts, 10 booked meetings, eight held meetings, six accepted discussions, three scoped proposals and one win. Positive-account rate is 5%; held rate is 80%; acceptance is 75%. These are invented planning inputs.

Economic input	Fictional value	Result
Retainer	USD 3,000 / month	Revenue per active client
Direct delivery cost	USD 1,800 / month	USD 1,200 monthly contribution
Cohort acquisition cost	USD 3,000	Research, systems, response and sales effort
New clients	1	USD 3,000 acquisition cost per client
Simple steady-state payback	3,000 / 1,200	2.5 months

If no client wins, acquisition cost per client is undefined, not zero. Payback also needs a positive contribution. Add onboarding costs and ramp when planning cash. Do not invent lifetime value from a short cohort.

Before increasing appointments, check delivery capacity. Ten good meetings can be harmful if the agency has no engineering, editorial or account-management capacity for the work it sells.

20. Test one uncertainty at a time

First decide whether the uncertainty concerns the account cell, first offer, evidence, contact role, copy or response handling. Test the earliest unresolved constraint. A subject-line test will not answer whether the agency is targeting the wrong business model.

Randomly allocate comparable accounts where practical. Keep all contacts from one account in the same test group. Hold research standards, contact frequency and response handling steady. Choose one primary outcome, such as accepted held discussions per eligible account, and retain the other funnel stages to explain the result.

Fictional test: offer A produces six accepted discussions from 200 accounts; offer B produces nine from 200. The observed difference is 1.5 percentage points. Under a simple independent-account normal approximation, the standard error is about 1.90 points and the approximate 95% interval is -2.22 to +5.22 points. That result is inconclusive; it does not establish that B wins.

Rare outcomes need more evidence or a more cautious decision. Report the allocation rule, sample size, period, maturity and exclusions. Treat repeated peeking as a limitation. If groups differ in firm size or intent, describe the comparison rather than calling it a controlled test.

21. Troubleshoot the pipeline at the failing stage

Observed pattern	Likely questions	First repair
Many wrong-person replies	Is the role mapped to the proposed decision?	Rebuild role and ownership rules
Assets requested, few discussions	Does the asset expose a real unresolved decision?	Improve usefulness and next-step clarity
Booked, not held	Were time, timezone, purpose and attendees agreed?	Repair scheduling and reminders
Held, not accepted	Is the offer attracting unsupported work?	Tighten account cell and meeting rules
Accepted, few proposals	Is implementation or value unclear?	Improve discovery and scoped assessment
Proposals, few wins	Is the economic case, proof or decision route weak?	Review decided deals and objections
Wins, weak contribution	Did scope, revisions or support exceed assumptions?	Repair pricing and delivery boundaries

Do not infer the cause from one week with a tiny denominator. Review actual messages, evidence cards, call notes and delivery estimates. Use one owner for each corrective action and a review date. Keep the original cohort visible while new campaigns run.

22. Run the first 90 days with decision gates

Period	Deliverable	Gate to proceed
Days 1-10	Account cell, 25 reviewed records, first-asset example, meeting rule	Research and delivery owners agree the scope
Days 11-30	Bounded outreach, reply coding, actual asset delivery	Review relevance, objections and response capacity
Days 31-60	Accepted discussions, paid assessments or scoped proposals	Confirm implementation route and credible economics
Days 61-90	Decided cohorts and first delivery reviews	Increase activity only where contribution and capacity support it

The dates are planning intervals. A long enterprise buying cycle may not produce a mature win cohort in 90 days. Record that honestly. Early success is improved evidence, appropriate discussions and a clearer decision path; realised commercial results require their own observation window.

23. SEO campaign message bank

These are reusable examples, not messages already sent. Replace all brackets with verified facts. The first four are initial offers for different cells. The next four form alternative follow-ups; choose only those relevant to the conversation. The final four handle real replies. Preserve one request per message and keep the final version under 100 words including subject, signature and required footer.

1. Ecommerce category review

Subject: [category] search pages

Hi [first name],

Your [category URL] groups products by [verified attribute]. Which combinations deserve their own search landing page?

We can review five category and filter examples, then map useful pages, questions and implementation owners.

Would that one-page sample help?

Thanks,
Shanay
[Business postal address]
Prefer no further emails? Reply stop.

2. Local enquiry ownership

Subject: the [area] enquiry path

Hi [first name],

Your location directory lists [verified area]. Who owns the local page, Business Profile and follow-up when an enquiry arrives?

We can map those responsibilities for five agreed locations before discussing a wider SEO programme.

Would you like the worksheet?

Thanks,
Shanay
[Business postal address]
Prefer no further emails? Reply stop.

3. B2B commercial measurement

Subject: from [use case] search to pipeline

Hi [first name],

Your [page] explains [verified use case]. How does the team connect those visits to accepted sales opportunities?

We can map the page, enquiry and qualification events before suggesting more content.

Would a one-page measurement brief be useful?

Thanks,
Shanay
[Business postal address]
Prefer no further emails? Reply stop.

4. Verified migration trigger

Subject: the [platform] launch checks

Hi [first name],

Your [dated announcement] confirms the move to [platform]. Is search monitoring part of the release plan?

We can map URL changes, launch checks and owners for one page family. The brief would show what still needs verification.

Would that help the launch review?

Thanks,
Shanay
[Business postal address]
Prefer no further emails? Reply stop.

5. Add scope, not pressure

Subject: Re: [original subject]

Hi [first name],

The sample would cover [five agreed examples], not a full-site audit. Each row would show the observation, possible business effect and next check.

Your team could use it without a meeting.

Would you like that version?

Thanks,
Shanay
[Business postal address]
Prefer no further emails? Reply stop.

6. Explain the measurement limit

Subject: Re: [original subject]

Hi [first name],

Public pages can show the customer path. They cannot reveal your qualified enquiries or revenue.

For [workflow], the brief would separate visible facts from the first-party checks needed before an SEO estimate.

Would that distinction help your review?

Thanks,
Shanay
[Business postal address]
Prefer no further emails? Reply stop.

7. Find an appropriate owner

Subject: who owns [page family]?

Hi [first name],

For [verified page family], the useful review needs someone who owns both the commercial goal and the implementation route.

We can start with one page sample and list the open decisions.

Are you the right person for that review?

Thanks,
Shanay
[Business postal address]
Prefer no further emails? Reply stop.

8. Close the loop

Subject: pausing the [workflow] review

Hi [first name],

I will pause the [specific asset] offer here.

If [workflow] becomes part of a later review, we can start with the source, owner and measurement questions rather than a full programme.

Would a later review date be useful?

Thanks,
Shanay
[Business postal address]
Prefer no further emails? Reply stop.

9. Deliver a requested resource

Subject: your [worksheet]

Hi [first name],

Here is the [link] you requested. Start with [specific field] for [workflow you named].

The sheet separates the public observation from the checks that need your team's data.

Would help with that field be useful?

Thanks,
Shanay
[Business postal address]
Prefer no further emails? Reply stop.

10. Existing agency reply

Subject: Re: [workflow]

Hi [first name],

Understood. The [specific review] could support your current team if that boundary is still open.

We would keep the work to [bounded deliverable], with your agency retaining the wider programme.

Would that narrow scope be useful?

Thanks,
Shanay
[Business postal address]
Prefer no further emails? Reply stop.

11. Buyer asks for pricing

Subject: Re: [workflow] pricing

Hi [first name],

The price depends on [scope driver] and who implements the changes.

We can quote the [bounded assessment] first, with its inputs, output and exclusions clear. A wider retainer would need the implementation review.

Would that assessment scope help?

Thanks,
Shanay
[Business postal address]
Prefer no further emails? Reply stop.

12. Buyer requests a discussion

Subject: the [workflow] review

Hi [first name],

Happy to review [question you raised]. We can inspect the sample, confirm the owner and decide whether a focused assessment makes sense.

Would [day, time, timezone] or [day, time, timezone] work?

Thanks,
Shanay
[Business postal address]
Prefer no further emails? Reply stop.

24. Working templates

Account and evidence record

Account ID; buying entity; country; commercial offer; page family; relevant role; source URL; source period; research date; exact observation; sample rule; inference; unknowns; first asset; delivery capability; permitted route; suppression; owner; next review.

Five-page sample

Selection rule; five URLs; page purpose; customer action; observed condition; capture context; evidence reference; possible consequence; alternative explanation; next check; implementation owner; explicit exclusions.

Search measurement dictionary

Property; search type; country/device; period; page cohort; brand rule and version; property clicks; represented query clicks; unallocated difference; enquiries; accepted leads; opportunity/win maturity; attribution limits; data owner.

Migration review

Release date and confidence; old/new URL inventory; redirects; canonical and internal-link checks; assets; analytics continuity; staging/production boundary; release owner; rollback plan; monitoring cohort; unresolved dependencies; sign-off.

Meeting handoff

Buyer request in their words; agreed purpose; contact role; workflow; source facts; unknowns; promised asset and delivery status; date/time/timezone; attendees; acceptance rule; next owner; decision and review date.

Assessment and proposal

Business decision; confirmed baseline; evidence limitations; deliverables; implementation responsibilities; exclusions; required access; timetable; fees; direct delivery assumptions; reporting definition; success and stop conditions; procurement route.

Experiment record

Uncertainty; account allocation rule; group sizes; offer difference; common controls; primary metric; maturity window; exclusions; results by funnel stage; uncertainty interval or explicit limit; decision; follow-up owner.

Monthly acquisition review

Eligible accounts; positive accounts; appointments due; held; accepted; opportunities; decided proposals; wins; acquisition cost; direct contribution; ramp; capacity; rejection reasons; one corrective action; owner and review date.

25. Sending and measurement operating notes

Use a genuine sender identity, appropriate contact route and a shared suppression list. Verify the rules for the actual jurisdiction and recipient type. US CAN-SPAM covers commercial B2B email; UK guidance distinguishes corporate subscribers from sole traders and other individual subscribers. These are different frameworks, not worldwide permission. C02 C03

Follow the relevant mailbox-provider authentication, identity and unsubscribe requirements. Gmail's requirements differ by sender class. Infrastructure cannot make an irrelevant message wanted. Do not use domain or inbox rotation to evade a stop instruction or provider enforcement. C01

Use replies and downstream stages to evaluate the programme. Open events are affected by privacy features and should not drive pressure-based follow-up. C05

For benchmark discipline, Belkins' current study reports 34,393 replies from 7,530,489 emails sent in 2025, excluding auto-replies and bounce notifications. That is about 0.457% replies per sent email. It is not a positive-account or accepted-meeting benchmark for SEO agencies. C04

26. A campaign is ready when these decisions are concrete

The reader should be able to name the account cell, explain the first asset, reproduce the evidence sample, identify the response owner and state what counts as an accepted meeting. The agency should also know who can deliver the work and what a client contributes after direct cost.

The resource's public examples and studies support those decisions. They do not establish a universal reply rate, a guaranteed ranking gain or a hidden list of active buyers. Use first-party evidence as it becomes available, update the source period and keep the commercial claim proportional to what the data can support.

Sources, periods and interpretation limits

Research checked 30 October 2026. Primary-source documents, original research publishers and first-party company pages are distinguished from fictional worked examples. A source date is not a current purchase signal.

S01. Pew Google AI-summary browsing analysis

22 July 2025; March browsing, April result reconstruction. Traditional-result clicks: 8% with AI summary versus 15% without.

Method: 900 US adults; 68,879 unique Google searches; 12,593 produced summaries in reconstructed results.

Limit: Observational. Search results collected April 7-17 differ in time from March browsing. Not a site-level causal uplift or loss.

Read the source.

S02. Ahrefs AI Overview CTR update

4 February 2026; December 2023 versus December 2025. AI Overview presence associated with about 58% lower top-ranking desktop CTR versus modelled counterfactual.

Method: 300,000 keywords: 150,000 with AIO and 150,000 informational without; aggregated GSC desktop CTR.

Limit: Not a universal 58% traffic loss. The estimate adjusts for a broader CTR decline; it is not the raw 7.3%-to-1.6% change.

Read the source.

S03. SparkToro zero-click search study

8 June 2026; January-April 2026. Reports 68.01% zero-click Google searches in its US browser-based analysis.

Method: Similarweb mobile and desktop clickstream; earlier studies used other panels.

Limit: Panel changes limit clean year-to-year causal interpretation. A zero-click search is not necessarily a failed business interaction.

Read the source.

S04. Google AI features and websites

Current documentation. Existing SEO foundations apply; Google names no additional technical requirement or special markup for AI Overviews and AI Mode.

Method: Official Google Search documentation.

Limit: Specific to Google features. Eligibility does not guarantee indexing, display, citation or traffic.

Read the source.

S05. Google SEO Starter Guide

Current documentation. Discovery, content, links and technical visibility matter; changes take variable time and ranking is not guaranteed.

Method: Official Google Search documentation.

Limit: Not a performance forecast or fixed ranking timetable.

Read the source.

S10. Google Search spam policies

Current documentation. Policies cover scaled-content, link, doorway and other abuse.

Method: Official policy documentation.

Limit: Do not infer a violation from a superficial audit or promise manipulative shortcuts.

Read the source.

S06. Search Console performance overview

Current documentation. Defines clicks, impressions, CTR, average position and dimension-specific aggregation.

Method: Official measurement documentation.

Limit: Property and page aggregation differ. Exported unavailable values may appear as zeros.

Read the source.

S07. Search Analytics API query

Current documentation. API returns top data subject to internal limitations and does not guarantee every row.

Method: Official API specification.

Limit: Paging does not establish complete query coverage; reconcile coverage and aggregation.

Read the source.

S08. Google site moves with URL changes

Current documentation. Migration planning includes URL mapping, redirects, verification and monitoring.

Method: Official operational documentation.

Limit: A public redesign is not proof of an SEO failure or an open migration project.

Read the source.

S09. Google local ranking guidance

Current documentation. Local results depend on relevance, distance and prominence.

Method: Official Business Profile documentation.

Limit: One search from one place is not a whole-market local ranking audit.

Read the source.

S11. Promethean digital agency industry report

2026 landscape; 2025 financial results. Reports average agency net margin of 13% in 2025.

Method: Report combines surveys of agency leaders and employees with observational industry data; bases vary by measure.

Limit: Broad digital-agency context, not SEO-only profitability or contribution margin. Credit: Promethean Research, Digital Agency Industry Report, prometheanresearch.com/digital-agency-industry-report.

Read the source.

C02. FTC CAN-SPAM compliance guide

Current guidance. Commercial B2B email is covered; accurate identity, postal address and opt-out duties apply.

Method: Official US guidance.

Limit: US rule, not a worldwide sending permission. Verify jurisdiction and all applicable obligations.

Read the source.

C03. ICO business-to-business marketing

Current guidance. UK rules distinguish corporate from individual subscribers; personal-data duties can still apply.

Method: Official UK guidance.

Limit: Corporate status does not erase all data or channel obligations. Sole traders differ.

Read the source.

C01. Gmail sender guidelines

Current documentation. Authentication, sender identity, unsubscribe and spam-rate requirements vary by sender class.

Method: Official provider requirements.

Limit: Authentication is not permission to contact or a guarantee of inbox placement.

Read the source.

C05. Apple Mail Privacy Protection

Current documentation. Mail privacy features affect what senders can learn about email activity.

Method: Official product documentation.

Limit: Open tracking is not a reliable meeting-quality measure; no numerical open-rate adjustment is assumed.

Read the source.

C04. Belkins cold-email response-rate study

2026 article, January-December 2025 campaign data. 7,530,489 sent emails and 34,393 replies; excludes auto-replies and bounce notifications.

Method: Vendor-owned client campaign dataset; unique replies divided by emails sent.

Limit: Not positive-account, booked, held or accepted-meeting rate. Not specific to these three agency sectors.

Read the source.

Public workflow source register

The following pages support the named workflow examples. None establishes an agency relationship or purchase intent.

- SO1. Wise: Converter combines currency information and a money-transfer action. Primary page. Checked 30 October 2026.
- SO2. Zapier: Directory organises app and integration discovery with product routes. Primary page. Checked 30 October 2026.
- SO3. Ruggable: Category navigation distinguishes rooms, sizes, materials and collections. Primary page. Checked 30 October 2026.
- SO4. The Grounds Guys: Directory lists local service areas and locations. Primary page. Checked 30 October 2026.
- SO5. Trainline: Journey pages and a route-search interface support ticket discovery. Primary page. Checked 30 October 2026.
- SO6. Atlassian: Migration content addresses movement to cloud products. Primary page. Checked 30 October 2026.

Method and provenance

LS Cold Email informed the relevance, first-offer, message and appointment-setting design. Its reference examples are not presented as LeadScorer campaign results. Published numerical claims retain their population, period and limits. Fictional examples demonstrate calculations and operating decisions; they do not imply observed clients, sector benchmarks or forecast outcomes.