

Google Ads for B2B SaaS: An Economically Sound 90-Day Pilot

Connect search intent, conversion goals and CRM feedback within a bounded learning budget.



HOW TO USE THIS RESOURCE

Google Ads for B2B SaaS: An Economically Sound 90-Day Pilot

SaaS marketers and founders starting search advertising or reassessing an existing channel with unclear commercial results.

Your practical outputs

- A prioritised search-intent and exclusion map
- An economic pilot budget with a downside case
- A tested measurement chain through CRM qualification
- A 90-day decision plan for continuation or stopping

Read, choose, apply

Start with the overview and decision rules. Choose a concrete constraint in your business. Work through the relevant templates using your own evidence, record open assumptions and agree a review with the accountable people. The complete online edition provides linked sources and contextual navigation.

Editorial context

Explicitly fictional cases and example figures are working aids, not client results or universal benchmarks. Frameworks structure decisions; they are not scientifically validated scores. Platform capabilities may change. Sources and review context are disclosed.

Read online: [Google Ads for B2B SaaS: An Economically Sound 90-Day Pilot](#)

CONTENTS

1. Choose the question a search pilot can answer	4
2. Build a search-intent map and exclusion policy	4
3. Connect the ad promise to an evaluable landing page	5
4. Choose conversion goals without confusing proxies and value	5
5. Allocate budget to interpretable experiments	6
6. Evaluate cohorts after the click	6
7. Define search-specific stage evidence	7
8. Derive bids and budget from commercial constraints	8
9. Test the click-to-CRM feedback loop	9
10. Calculate a fictional search pilot	9
11. Diagnose search failures in the right order	10
12. Make search and sales share a usable feedback packet	11
13. Stage the 90-day search experiment	12
14. Copy the search pilot decision workbook	12
Turn this resource into a team brief	15
Sources & further documentation	16

1. Choose the question a search pilot can answer

Search advertising can test whether existing searches contain reachable, commercially relevant demand for a supported use case. It cannot establish the full size of a future market or prove that all attributed sales are incremental. Brand queries, for example, may capture people who already decided to visit. Keep that distinction visible in the pilot brief.

Use the original INTENT contract: Issue, Need, Task, Exclusions, Next step, Test. Describe the business issue behind a search, the likely need, the task the buyer wants to complete, unsuitable interpretations, the offered next step and the test that would support investment. Write contracts for a few coherent themes before creating a large keyword list.

A useful initial question is whether operations managers searching for approval-workflow software will accept a relevant evaluation conversation. A weak question is whether Google Ads can produce more traffic. The first links intent to a commercial outcome; the second rewards activity without defining quality.

Keep brand and non-brand evidence separate. Keep competitor comparison evidence separate if used, and obtain appropriate review of claims and trademark-sensitive wording. A small pilot benefits from interpretable boundaries more than from maximum coverage.

2. Build a search-intent map and exclusion policy

Group searches by intended task: finding a software category, solving a named workflow problem, comparing approaches or learning a concept. Avoid assuming that a phrase containing “software” always signals a purchase project. Inspect likely interpretations with sales and product specialists, then design the landing promise accordingly.

Google provides a search terms report for reviewing searches that triggered ads. Its official documentation also explains using irrelevant terms to inform negative keywords. Treat the report as observed evidence, not a complete census of every buyer thought. Review actual visible terms and downstream inquiries together; a term that looks relevant can still attract an unsupported use case.

Maintain three exclusion layers: categorically irrelevant meanings, commercially unsupported use cases and pilot-only boundaries such as an unserved geography. Record why each exclusion exists. A word such as “free” may describe unsuitable intent in one context but a valid comparison question in another. Blind blanket lists can suppress useful demand.

The negative-keyword review owner should check collateral effects before applying a broad change. Keep a dated log of the change and inspect whether the intended query class disappeared. The objective is a cleaner test, not the longest possible negative list.

3. Connect the ad promise to an evaluable landing page

Give each coherent search theme a page that answers the likely buying question. State the supported workflow, suitable company context, relevant capabilities, constraints and the offered next step. A generic homepage may make a visitor reconstruct this argument alone. A highly specific page should still accurately represent the product.

Draft the promise in plain language before adapting it to available ad fields. Check that every material claim can be fulfilled. “See an example approval workflow” is easier to verify than “Transform operations instantly”. Avoid unsupported superiority claims and invented numerical improvements. A credible limit can improve fit by helping unsuitable buyers self-select.

Use one primary next step and an accessible secondary learning path when appropriate. A technical evaluator may need documentation before requesting a meeting. State whether submitting a form requests a demo, an assessment or only the resource. Do not disguise a sales request as an unconditional download.

Review the page on a phone, with slow loading and with a failed form submission. Confirm that validation does not erase entered information and that success is recorded once. Technical reliability is a commercial requirement: paying for a visit to a broken or misleading page cannot answer the pilot question.

4. Choose conversion goals without confusing proxies and value

Define distinct events for form submission, qualified inquiry, opportunity and closed outcome. Google's documentation distinguishes qualified and converted leads and describes enhanced conversions for leads as a way to supplement offline conversion matching with user-provided data. Verify eligibility, data handling and actual implementation in the account; a feature description is not proof your import is correct.

Choose the optimisation event according to meaningfulness, reliability, timeliness and available signal. A deep event may be valuable but too delayed or sparse to guide the initial pilot. A shallow event may be frequent but commercially weak. Record the compromise explicitly and retain downstream quality checks. Do not pretend every submitted form has the same economic value.

When importing later outcomes, preserve the event time and a stable business identifier, follow the current vendor specification and test correction handling. If the same event is sent through more than one path, implement the supported duplicate-control approach. Reconciliation with CRM remains necessary.

Do not switch goals repeatedly to make the latest report look better. Change them when evidence shows the existing signal is broken or misleading, annotate the date and evaluate later cohorts separately. Platform optimisation cannot repair a fundamentally unsuitable offer.

5. Allocate budget to interpretable experiments

Set a hard cash limit for the entire pilot and a smaller initial release. Reserve room for measurement repair and one meaningful correction. Dividing a small budget across many markets, products and experiment variants can leave every result too thin to interpret. Choose one supported segment and a limited set of search themes.

Derive a click-cost ceiling from economics. If the assumed allowable qualified-lead cost is €500 and the assumed click-to-qualified-lead probability is 3%, the implied allowable CPC is €15. Both numbers are fictional planning assumptions. If actual commercial evidence is worse, the ceiling falls; a high click-through rate does not override it.

Record the current bidding approach, campaign boundaries and reason for each choice. Use official in-account guidance for available settings and requirements rather than treating a particular bidding strategy as universally correct. Change one major lever at a time where practical and note the dates.

Check pacing and data integrity frequently, then make commercial decisions on an agreed review schedule. A single expensive click should not trigger a strategic rewrite. A broken conversion event, irrelevant query class or exhausted sales capacity does require immediate action. Distinguish emergency corrections from hypothesis changes in the log.

6. Evaluate cohorts after the click

The search team needs feedback from actual conversations. Give sales a concise view of the inquiry context and ask for usable reasons when rejecting or advancing it. A generic quality score without an explanation is difficult to translate into a keyword, landing-page or offer decision.

Analyse acquisition cohorts by meaningful intent group and date. Track the number of valid inquiries, qualified leads, accepted opportunities and eventual customers. Include unresolved counts and age. Do not compare last month's new leads with last quarter's fully matured wins as though their conversion windows were equivalent.

When performance looks strong, ask whether brand capture, duplicate people, existing customers or one unusually large opportunity explain the result. When performance looks weak, ask whether data loss, missed follow-up or a long buying cycle explains the apparent gap. These checks protect against both false optimism and premature shutdown.

The pilot ends with a decision memo: what intent was tested; what was learned; what remains unobserved; which costs were included; and whether to continue, modify, pause or stop. Search

can be commercially useful without becoming the largest channel. The question is whether the next bounded investment has a defensible role in the revenue system.

Technical capabilities and sources

The search terms report supports inspection of observed searches associated with ads. [About the search terms report](#)

Google describes using the search terms report to identify negative keyword candidates. [Use search terms for negative keyword ideas](#)

Enhanced conversions for leads supplements offline conversion matching with user-provided data; verify data requirements in the linked guidance. [About enhanced conversions for leads](#)

Google documents qualified-lead and converted-lead conversion categories. [About qualified leads and converted leads](#)

GA4 campaign URL parameters identify referring campaigns in acquisition reporting. [Campaign URL parameters in GA4](#)

The GA4 event reference documents supported event structures; use the current specification when implementing events. [GA4 Measurement Protocol events reference](#)

Lifecycle stages can categorise contact and company progress; configure the organisation's acceptance rules separately. [Use HubSpot contact and company lifecycle stages](#)

CRM lead scores can use actions and properties; a score remains an aid to prioritisation rather than proof of a purchase project. [Understand HubSpot lead scoring](#)

7. Define search-specific stage evidence

Create a lead definition that can distinguish a useful search inquiry from a successful tag firing. The minimum record contains a usable contact, a submitted request, the campaign context available under the approved measurement design and an eligibility decision. Keep spam, internal tests, support requests and employment inquiries outside the acquisition denominator under documented rules. Retain their counts as operational diagnostics.

Qualification should answer whether the company can use the product and whether its request provides a relevant business reason to continue. The SDR records that reason in ordinary language. "Searching for workflow software" explains how someone arrived, but not whether there is a supported project. Sales creates an opportunity only when discovery establishes an actual buying process and a next action. An automated calendar booking alone need not meet that rule.

Store first submission, qualification, opportunity creation and outcome dates separately. A qualification correction should not overwrite the original acquisition date. When an existing contact returns through a paid query, preserve the new interaction while following the agreed rule for unique leads. Otherwise remarketing and repeat evaluation can inflate acquisition.

Check ten sample inquiries with the campaign manager and SDR before interpreting a conversion column. The two people should identify the same eligible records and explain every exclusion.

Search object	Evidence	Exclusion
Valid inquiry	Usable request, recorded source	Tests, spam, support
Qualified inquiry	Supported workflow and business reason	No relevant need
Search opportunity	Sales acceptance and next action	Booking without discovery

8. Derive bids and budget from commercial constraints

Start the pilot spreadsheet with the maximum cash the business can lose while learning. Next estimate an affordable media acquisition cost after sales and onboarding allocations. Work backwards through the expected funnel, making every probability an explicit assumption. Affordable CPC equals allowable media CAC multiplied by click-to-customer probability. The equivalent staged expression multiplies click-to-valid-lead, qualification, opportunity and win probabilities; it is useful only when stages and denominators are compatible.

Assume an allowable media CAC of €5,000, 5% click-to-valid-lead rate, 30% qualification, one-third qualified-to-opportunity progression and 20% opportunity win rate. These fictional assumptions imply a 0.1% click-to-customer probability and a €5 affordable CPC. Doubling the lead conversion rate would double that modelled ceiling only if downstream quality stayed constant. That condition is easy to violate by weakening the form or broadening the promise.

Model a downside before releasing spend. Reduce the least-supported probability, extend the sales cycle and include creative or implementation work. If a plausible downside would exceed available cash, narrow the experiment or do not launch. Competitor CPC estimates and platform forecasts can inform planning, but neither can establish your conversion quality or acceptable acquisition cost. Review actual costs against the written envelope instead of retroactively raising the ceiling to justify expenditure.

Planning input	Fictional assumption	What to inspect
Allowable media CAC	€5,000	Finance-approved cost scope
Click-to-customer	0.1%	Compatible stage probabilities
Affordable CPC	€5	Downside if quality falls

9. Test the click-to-CRM feedback loop

Draw the data path on one page: advertising interaction, landing request, form acceptance, CRM record, sales classification and later conversion feedback where implemented. Identify which system owns each timestamp and what key connects the steps. Record a safe recovery procedure for temporary integration failure. Data sent twice during a retry should not create two business outcomes.

Use a marked internal inquiry to test the whole route. Verify campaign context at the form, successful record creation, assignment to the intended owner, a qualification decision and the downstream import result. Then submit the same person again and inspect duplicate behaviour. Repeat with consent choices that should restrict tracking; the site must respect them while fulfilment behaves according to the approved design. Do not treat a tag manager preview as proof that the CRM import works.

For enhanced conversions for leads, follow the current Google specification for supported user-provided data, matching and import requirements. Hashing does not remove the need for appropriate data handling. Confirm accepted events and error messages in the actual account. Compare imported qualification counts with eligible CRM events, allowing for known processing delay and matching limits. Maintain an exception queue for rejected or missing imports. Only the implementation owner should close an exception after a tested correction, with the campaign manager informed of any reporting break.

Import test	Expected result	Owner
Repeated submission	No duplicate business event	CRM specialist
Qualified stage	Correct time and mapped action	RevOps
Restricted tracking choice	Approved behaviour respected	Implementation owner

10. Calculate a fictional search pilot

FICTIONAL SEARCH PILOT — Assume Nordlicht Workflow spends €12,000 on media for 90 days. At an assumed €8 CPC it receives 1,500 clicks. Assume 75 valid inquiries, 30 qualified inquiries, ten opportunities and two eventual customers. All figures are invented assumptions. The resulting media CPL is €160, cost per qualified inquiry €400, cost per opportunity €1,200 and media CAC €6,000. The 90-day activity window does not imply that all wins are observable by day 90.

Assume each customer contracts €12,000 recurring annual value with 80% gross margin. Two wins produce €24,000 contracted ARR, while ten opportunities represent €120,000 unweighted first-year pipeline. Neither amount is recognised accounting revenue. Add €8,000 sales and implementation cost allocated to acquisition, and fully loaded CAC becomes €10,000. At €800

monthly gross profit per customer, simple payback becomes 12.5 months under constant revenue and margin assumptions.

Now inspect the intent groups. Assume a workflow-specific group produced eight of the ten opportunities, while a broad research group produced two and most rejected inquiries. This is a hypothesis for a bounded allocation change, not proof of a universal keyword rule. Keep existing outcomes attached to their original cohorts. Narrow the next test and check whether commercial quality survives rather than projecting the best group's rate across unlimited spend.

Search pilot measure	Calculation	Fictional value
Media CPL	€12,000 / 75	€160
Media opportunity cost	€12,000 / 10	€1,200
Loaded CAC	€20,000 / 2	€10,000

11. Diagnose search failures in the right order

First inspect delivery and implementation: correct geography, working destination, accurate promise and successful submission. Next inspect visible search terms for unsuitable meanings. Then inspect accepted and rejected inquiries. Only after these checks should the team debate creative refinements or more aggressive bidding. A tracking outage can make a healthy group look weak; a duplicated success event can make an irrelevant group look attractive.

Separate three patterns. Relevant searches with few submissions suggest a page, offer or technical problem. Many submissions with poor qualification suggest an intent or promise mismatch. Good qualified conversations with few opportunities suggest discovery, timing, capability or pricing barriers that require sales or product evidence. These are starting hypotheses, not automatic diagnoses.

Use a query-review worksheet with term, intended task, observed inquiry quality, proposed action and collateral-risk check. A single rejected inquiry is not enough to exclude a broad useful concept. Repeated requests for a capability the product cannot deliver justify revising the promise or exclusion policy. Stop the affected group immediately if it systematically advertises unsupported functionality. Pause expansion when classification feedback is unavailable. Without that feedback, the team cannot tell whether lower form cost is genuine improvement or cheaper irrelevance. Reopen the test only after the feedback owner and verification date are recorded.

Search symptom	Inspect first	Bounded action
Irrelevant inquiries	Visible terms and actual requests	Exclude unsuitable meaning
Relevant visits, no requests	Promise, form and destination	Repair one page constraint

Search symptom	Inspect first	Bounded action
Imports missing	CRM-to-platform exception queue	Pause expansion and repair

12. Make search and sales share a usable feedback packet

The campaign manager needs an actionable reason, not a weekly complaint that leads are weak. Agree a compact classification packet: inquiry reference, source context, supported-use-case status, contact outcome, qualification reason, opportunity link and next action. Avoid requiring sales to complete a lengthy marketing questionnaire before responding to the buyer. Necessary fields should serve both operational follow-up and analysis.

An SDR returns unsupported applications to marketing for promise review, ambiguous cases for clarification and valid but premature projects to the approved later-contact path. A missed response is an operations exception, not a targeting failure. Sales leadership owns acceptance consistency; the search specialist owns query and campaign changes. The CRM owner owns reliable linkage and correction handling. Finance owns the included-cost definition.

Set a regular feedback cut-off so the campaign manager knows which recent records are still unclassified. Record staffed response windows and substitute owners. When volume exceeds capacity, reduce acquisition or change the fulfilment promise rather than silently letting inquiries age. Test the packet during a real review: can the specialist decide whether to inspect a query, a page or a routing rule? If not, replace vague categories with a short factual explanation. Close the loop by telling sales which campaign change followed its evidence and what later outcome will be reviewed.

Feedback participant	Provides	Owens next action
SDR	Fit and request evidence	Classification by review cut-off
Search specialist	Term and promise context	Dated campaign correction
CRM owner	Linkage and import status	Verified repair and replay

13. Stage the 90-day search experiment

Days 1–30 establish the experiment boundaries. Approve one segment, a small set of intent contracts, the offer, exclusions, cash limit and conversion specification. Build and test the destination and CRM route. Release the first budget tranche only after a valid inquiry can be followed through classification. Document the actual account settings rather than relying on screenshots from an older interface.

Days 31–60 focus on query relevance and commercial feedback. Review pacing, implementation errors and rejection patterns. Make one meaningful correction with a recorded rationale. If a query group attracts the wrong use case, repair the promise or exclude the unsuitable meaning. If valid requests go unanswered, fix ownership before buying more traffic. Preserve pre-change and post-change acquisition dates for later cohort comparison.

Days 61–90 produce a provisional commercial decision. Compare the intent groups at equivalent ages and show how many opportunities remain unresolved. Approve a bounded continuation only if the quality evidence, cost envelope and sales capacity support it. A long sales cycle may require stopping media while observing existing opportunities. Scaling is a new decision with a new cash limit, not the automatic next stage of a pilot. Assign a date for a final outcome update and specify what would overturn the provisional conclusion.

14. Copy the search pilot decision workbook

Complete this worksheet before the next budget release. Pilot question: __; supported segment: __; intent group: __; unsuitable meanings: __; promised next step: __. Media cash limit: __; non-media acquisition allocation: __; allowable CAC: __; assumed click-to-customer probability: __; implied affordable CPC: __. Record the downside assumptions alongside the base case.

The weekly inspection sheet needs campaign dates, observed terms, valid inquiries, classified inquiries, qualified inquiries, opportunities, import exceptions and unworked records. Attach a few factual examples behind the largest rejection category. Keep personal information in the authorised CRM and use references in the review document. “Unknown” should remain visible when a record has not been worked.

The decision memo asks four different questions. Did the intended searches reach the page? Did the promise produce relevant requests? Did sales establish plausible buying processes? Did mature outcomes fit the approved economics? A yes to the first is not a yes to the fourth. State which question the available evidence can answer.

Continue for a bounded period when quality is credible but outcomes are immature. Modify a specific repairable constraint. Pause broken measurement or follow-up. Stop a contradicted commercial hypothesis without an affordable correction. Name the owner, amount, review date and restart condition for the chosen action. Keep this memo with the account change log so later performance can be interpreted.

Audit the marginal expansion before releasing more budget

A successful initial group does not imply that the next tranche will perform identically. The first spend may concentrate on the most explicit queries, a favourable geography or a short period of unusually strong demand. Expansion can introduce broader meanings and different auction conditions. Before increasing the limit, describe what inventory or audience the additional money is expected to reach. If nobody can explain that change, the expansion decision is incomplete.

Create an expansion card: original intent boundary; proposed new boundary; additional cash; expected staff load; unchanged quality rule; risk to existing coverage; and rollback condition. Keep the initial cohort visible rather than blending it into the expanded campaign. If the new spend produces more valid inquiries but worse qualification, report both effects. Do not use average historical CAC to conceal a deterioration in the next increment.

The first expansion can be deliberately modest. Its purpose is to test whether useful economics survive a changed scale or scope. A conservative release is not a universal percentage rule: the amount depends on cash, observed volume, signal delay and available sales capacity. Assign one person to check pacing and another to confirm commercial classifications. If feedback slows as volume rises, treat that as a capacity constraint before concluding the advertising quality changed.

Close the pilot so another person can operate it

The closing package contains the approved intent contracts, actual settings, exclusion reasons, landing promise, event specification, import exception procedure, cohort results and decision memo. Include access ownership and the location of the current implementation instructions. Avoid leaving the only working knowledge in one specialist's personal notes.

Ask a substitute owner to trace one inquiry and explain one exclusion. Ask finance to reproduce the cost scope and sales to identify the next action for an open opportunity. Resolve any discrepancy before handover. Archive superseded assumptions with dates rather than deleting them. Finally, schedule the remaining cohort outcome review and name the person who will update the provisional decision after media has ended. The pilot is complete when its evidence can be reused and its open obligations remain owned.

Primary sources and review date

- [About the search terms report](#) — 2026-10-24
- [Use search terms for negative keyword ideas](#) — 2026-10-24
- [About enhanced conversions for leads](#) — 2026-10-24
- [About qualified leads and converted leads](#) — 2026-10-24
- [Campaign URL parameters in GA4](#) — 2026-10-24
- [GA4 Measurement Protocol events reference](#) — 2026-10-24
- [Use HubSpot contact and company lifecycle stages](#) — 2026-10-24

- [Understand HubSpot lead scoring](#) — 2026-10-24

YOUR TRANSFER CANVAS

Turn this resource into a team brief

Use this page for a joint review. Agree one practical output, observable evidence and an accountable owner. Add the specific methods and templates from the relevant chapters.

Working field	Your agreement
Priority output	A prioritised search-intent and exclusion map
Your starting point	Finding / source / date: _____
Smallest useful test	_____ _____
Evidence & counter-evidence	_____ _____
Ownership & review	Owner / reviewer / date: _____
Stop or change when ...	_____ _____
Decision & next action	_____ _____

Keep unresolved assumptions explicitly open. A completed worksheet is not a substitute for evidence or accountable judgement.

Sources & further documentation

Primary sources checked 24 October 2026. Sources support the linked technical or factual statements; the working models and fictional cases are original editorial decision aids. Recheck changing platform details before implementation.

1. [About the search terms report](#)
2. [Use search terms for negative keyword ideas](#)
3. [About enhanced conversions for leads](#)
4. [About qualified leads and converted leads](#)
5. [Campaign URL parameters in GA4](#)
6. [GA4 Measurement Protocol events reference](#)
7. [Use HubSpot contact and company lifecycle stages](#)
8. [Understand HubSpot lead scoring](#)

Next step

If you want help adapting the working plan to your product, evidence and team capacity, request a non-binding strategy conversation.

patrickwings.com