

B2B Demand Generation: A Measurable Pipeline Operating System

Connect buyer decisions, accountable handovers, reliable cohorts and economically bounded experiments.



HOW TO USE THIS RESOURCE

B2B Demand Generation: A Measurable Pipeline Operating System

SaaS founders, marketing leaders, revenue operations and sales leaders selling a considered B2B product.

Your practical outputs

- A documented target segment with exclusions
- An offer and buyer-decision map
- A measurable pipeline with defined handovers
- An economically bounded 90-day learning plan

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: [B2B Demand Generation: A Measurable Pipeline Operating System](#)

CONTENTS

1. Start with the constraint, not the channel list	4
2. Define a segment through a costly business problem	4
3. Map buyer decisions before producing content	5
4. Separate demand creation, capture and conversion	5
5. Design an offer ladder that respects readiness	6
6. Run an evidence meeting that changes work	6
7. Create a shared definition of progress	8
8. Set the economic envelope before scaling	9
9. Build a measurement chain that can be audited	10
10. Work through a fictional B2B SaaS case	11
11. Diagnose failure by locating the broken transition	12
12. Make ownership and handovers operational	13
13. Run a 30/60/90-day learning cycle	14
14. Use the workbook and make an evidence-based decision	14
Turn this resource into a team brief	17
Sources & further documentation	18

1. Start with the constraint, not the channel list

A pipeline operating system connects choices that are often managed separately: whom to serve, which problem to make understandable, which next step to offer, how records move and how commercial evidence changes investment. Its purpose is dependable learning and coordination. It cannot create product-market fit, shorten every procurement process or turn attention into purchase intent by definition.

Use the original CLEAR loop: Constraint, Learning question, Evidence, Accountable action, Review. Start with one constraint that prevents commercially relevant progress. “We need more demand” is too broad. “The right companies request demos, but most cannot explain a priority workflow” is testable. The learning question becomes whether a narrower assessment offer produces better problem evidence. Evidence must include conversation content, not just the number of requests.

Write one operating card before selecting a tactic. State the constrained transition, the affected segment, the suspected cause and the smallest change that could resolve it. Then name the owner who can act on the evidence. A marketing team should not claim to fix a product requirement it cannot change; it should escalate the constraint to the product owner.

CLEAR is a management framework, not a causal attribution model. Several constraints can coexist, and an apparently local issue may originate earlier. Revisit the constraint after each learning cycle rather than protecting the original hypothesis.

2. Define a segment through a costly business problem

Describe the target company using conditions that make the problem relevant: operating complexity, workflow frequency, existing systems, responsibility and the consequences of doing nothing. Industry and employee count are useful filters, but they are insufficient explanations of need. A hundred-person business with a simple process may be less suitable than a smaller company handling a regulated workflow.

Create a segment card with three layers: structural fit, problem evidence and reachable buying roles. Structural fit defines where the product can work. Problem evidence explains why change might matter now. Reachability determines whether a useful conversation is possible within the available channels and budget. Do not confuse reachability with fit; the easiest audience to advertise to may be commercially weak.

Interview a small set of customers, lost prospects and sales colleagues. Ask what happened before evaluation began, what workarounds existed and what would have made the project unnecessary. Record the actual phrases buyers use. Do not treat a handful of interviews as a prevalence estimate; they reveal hypotheses and language rather than representative market size.

Write exclusions with equal care. Unsupported integrations, unsuitable deployment requirements and unaffordable onboarding effort should appear in targeting or qualification rules where appropriate. A narrow test segment is reversible. An ambiguous promise that attracts companies you cannot serve consumes both acquisition budget and buyer trust.

3. Map buyer decisions before producing content

A content inventory is not a decision map. Organise the journey around questions the buying group must resolve: Is this problem worth addressing? Which approaches are credible? What does implementation require? What could go wrong? How will the business justify expenditure? A buyer may revisit these questions in a different order as new colleagues join.

For each decision, assign a useful asset, the evidence it contains and an appropriate next action. A diagnostic worksheet can help frame the problem; an implementation guide can expose dependencies; a security explanation can reduce technical uncertainty; a business-case template can make assumptions visible. A product demo belongs where seeing the workflow answers an actual evaluation question.

Distinguish evidence from assertion. “Easy implementation” is a claim. A documented list of required systems, tasks, owners and known limits is useful decision support. When using customer examples, secure approval and state context; never imply that one result is a universal expectation.

Build the first map with sales and customer success. Ask which unanswered questions repeatedly delay a next step. Prioritise one decision that can be improved with available expertise. The result should be a connected sequence of useful choices rather than an automatic march from article to download to sales call. Offer direct access when a buyer is already ready to evaluate.

4. Separate demand creation, capture and conversion

Demand creation makes a costly problem and a credible alternative understandable. Demand capture responds when people actively look for a category, solution or supplier. Conversion helps an interested buying group complete a meaningful next step. These functions overlap, but separating their purposes prevents one reporting rule from distorting all three.

An educational LinkedIn campaign may reach relevant roles without immediate inquiries. A search campaign may capture existing demand that another activity helped create. A conversion improvement may increase demo requests without expanding market demand. Assign a primary question to each activity and keep its downstream commercial evidence visible. Do not excuse indefinite unmeasurable spending by calling everything brand building.

Use a portfolio card for each activity: intended audience, buying question, distribution mechanism, expected observable signal, downstream quality check and review horizon. The horizon should reflect the type of evidence and actual buying cycle. It is not permission to ignore early evidence of poor fit or broken execution.

When budgets are small, connect one creation activity to one capture path and one conversion offer for the same segment. This creates a coherent system that can be reviewed. It still does not permit summing platform-attributed revenue, or assuming that a later search interaction proves the earlier impression caused a sale. Use buyer-reported discovery and carefully designed comparisons as additional evidence.

5. Design an offer ladder that respects readiness

An offer is a contract about the next useful step. State what the person receives, how much effort it requires, who it suits and what happens afterward. A guide offers independent learning; an assessment offers structured diagnosis; a demo shows product behaviour; a commercial conversation explores a purchase. These are different promises and should not share an undisclosed sales follow-up policy.

Build an offer ladder with three rungs: learn, diagnose, evaluate. Let buyers skip rungs. Someone with an approved project may reasonably want a technical evaluation immediately. Someone exploring a problem should not have to book a meeting to access a basic explanation. Offer friction should reflect the value and operational work involved, not a desire to maximise database size.

For each rung, write a completion criterion. Learning is complete when a person can frame the problem; diagnosis when relevant constraints are documented; evaluation when an agreed product question has been tested. These criteria are aids for service design, not proof of psychological readiness.

Audit the landing page against the actual fulfilment. If it promises a personalised assessment, the team must have the capacity and competence to provide one. If it offers a freely readable resource, do not imply that a gate or consultation is required. Later gating is a separate choice that should preserve clear expectations and appropriate permissions.

6. Run an evidence meeting that changes work

Hold a short weekly operating review around the constraint card. Start with exceptions: lost submissions, ignored requests, repeated misunderstandings and commercially unsuitable inquiries. Then review the relevant cohort and its stage age. A dashboard is useful as an index to evidence; the actual decision often requires a few records and a conversation with the owner.

Use four questions: What did buyers do? What do we know about why? What can we change? What would prove the change was useful? Keep observations separate from stories. “Seven rejected inquiries requested an unsupported integration” is an observation. “Our entire market wants that integration” is a hypothesis requiring broader evidence and a product decision.

Maintain a decision log with date, owner, evidence, alternative explanation and revisit condition. Close completed actions explicitly. Reopening the same issue every week without a decision usually indicates missing authority or an unclear commercial definition, rather than insufficient reporting.

The monthly meeting should reconcile marketing signals with sales capacity and finance outcomes. The quarterly meeting can reconsider the target segment and allocation. Do not change strategy in response to every weekly fluctuation, but do not protect a weak strategy from repeated contradictory evidence. The operating system works when it makes a decision easier to explain and an agreed action easier to complete.

Technical capabilities and sources

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](#)

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](#)

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](#)

LinkedIn conversion tracking requires a data source, defined conversions and association with ad sets. [Get started with LinkedIn Conversion Tracking](#)

HubSpot workflow settings include controls for enrollment and suppression, subject to the workflow type. [Manage HubSpot workflow settings](#)

7. Create a shared definition of progress

A demand programme becomes measurable when every transition has an observable rule. Use the following working definitions before discussing channel performance. A lead is a distinct, usable contact record associated with a relevant inquiry or content request. Qualification requires documented fit and a meaningful reason for a business conversation. An opportunity exists only after sales accepts a specific buying process, with a problem, an accountable contact, a next step and a plausible commercial scope. Revenue requires a separately defined financial event. A signed contract, annual recurring revenue and recognised accounting revenue are different measures.

Object	Entry evidence	Counting rule	Accountable owner
Lead	Valid contact and recorded action	Unique contact in cohort	Marketing operations
Qualified lead	Fit plus explicit business evidence	First qualification per agreed period	Marketing and SDR lead
Opportunity	Sales acceptance and agreed next step	Unique CRM opportunity	Sales
Customer / revenue	Signed agreement or finance event	Contract or accounting rule	Finance

Do not convert these definitions into a rigid universal funnel. An existing account can open a new opportunity without creating a new lead. Several people can contribute to one opportunity, and a recycled contact can return months later. Keep person, account and opportunity identifiers separate, then connect them with explicit relationships. Missing relationships belong in a data-quality queue rather than being silently estimated.

For the pipeline operating system, the decisive evidence is a relevant business problem, segment fit and an agreed commercial next step. A weak signal such as a page view can justify further observation, but it cannot replace that evidence. Ask sales to classify ten recent records independently, compare disagreements, and change the written rule until a second person can reproduce the decision. Keep the original timestamps when definitions change; add a definition version instead of rewriting history. This protects cohort comparisons and prevents an apparent improvement caused only by a looser qualification rule.

Copyable acceptance statement: “A record enters this stage when ___ is recorded by ___; duplicate handling is ___; exceptions are reviewed on ___.” Complete it for all four objects before launch.

8. Set the economic envelope before scaling

Start with the decision the business can afford, then work backwards. A channel budget is not justified by an attractive cost per lead alone. It must be compatible with the cost of converting leads, the time before contracts close, gross margin, onboarding effort and the cash the company can commit while learning. Use contribution assumptions agreed with finance rather than a revenue multiple copied from another SaaS company.

A simple planning model is: allowable acquisition cost = monthly recurring revenue per new customer × gross-margin rate × chosen payback months. This is a planning ceiling, not an accounting standard. If onboarding and sales costs are excluded, label the result a media-only allowance only after subtracting those costs. Expansion, churn and financing constraints may require a different ceiling. Do not call a lifetime-value estimate reliable when retention history is short or customer segments behave differently.

For an explicit example, assume €1,000 monthly recurring revenue, 80% gross margin and a 12-month payback limit. The gross-margin envelope is €9,600 per acquired customer. Assume €4,600 allocated to sales, setup and other acquisition work; the remaining media allowance is €5,000. If the assumed lead-to-customer probability is 4%, the implied allowable media CPL is €200. If that probability falls to 2%, it becomes €100. These are fictional inputs, not benchmarks or recommended targets.

Input	Team assumption	Downside case	Decision consequence
Media budget	Fill in €__	Same committed spend	Maximum cash at risk
Qualification rate	Fill in __%	Half the estimate	Recalculate cost per qualified lead
Opportunity win rate	Fill in __%	Lower than base case	Recalculate customer cost
Time to close	Fill in __ days	One extra cycle	Extend observation, not automatic spend

Apply this envelope to the pipeline operating system: allocate money across observed bottlenecks rather than giving every channel an equal share. Freeze expansion if the downside case cannot fit the approved learning budget. A positive model licenses an experiment; it does not demonstrate incremental demand or future profitability. Record both the base case and the reason it might be wrong.

9. Build a measurement chain that can be audited

Write the measurement specification before implementing tags. For each event, specify its business meaning, system of record, identifier, timestamp, currency, consent requirements, duplicate key and correction procedure. A successful form display is different from a successful submission, and a successful submission is different from a sales-accepted opportunity. The event names should reveal these differences to someone reading the report six months later.

Use campaign identifiers consistently across links, forms and CRM records. Keep a controlled dictionary for source, medium, campaign and creative values. Preserve the originally captured values alongside later corrections. The official GA4 documentation describes how campaign URL parameters identify referring campaigns; that technical capability does not guarantee complete attribution. Browser restrictions, consent choices, forwarding and offline interactions create gaps. Read the linked source before choosing implementation details.

For this programme, the critical join is the person-to-account-to-opportunity relationship across source capture and sales acceptance. Test it with a clearly marked internal record that moves through submission, assignment, qualification, opportunity creation and rejection. Check both a happy path and a failure path: duplicate submission, missing company, declined permission, reassigned owner and a later correction. A dashboard that works only for a perfect record is not ready for a paid launch.

Test	Expected evidence	Failure response	Sign-off
Valid submission	One record and one event	Pause release and inspect trigger	Operations
Duplicate	Existing person correctly linked	Fix merge or idempotency rule	CRM owner
Qualification	Correct timestamp and reason	Repair mapping; replay safely	Sales operations
Withdrawal / exclusion	Suppression applies as designed	Stop affected automation	Responsible owner

Reconcile three totals weekly: source submissions, CRM-created records and records eligible for reporting. Investigate the difference instead of forcing agreement. Platforms may credit the same commercial outcome under different attribution rules, so their reported conversions cannot simply be added together. Maintain one CRM opportunity count and one finance revenue count for management decisions. Attribution provides a useful lens on observed activity; it does not establish causality.

Release checklist: event specification approved; test records excluded from reporting; access limited to necessary staff; retry procedure documented; rollback owner reachable. Where a

platform feature or connector depends on a subscription, verify availability in the actual account before promising a workflow.

10. Work through a fictional B2B SaaS case

FICTIONAL WORKED EXAMPLE — Every number in this chapter is an explicit planning assumption for “Nordlicht Workflow”, an invented B2B SaaS company. Nothing here describes a client result, a market benchmark or a forecast. Assume an operations product sold to mid-sized businesses, €12,000 first-year recurring contract value, 80% gross margin and a buying cycle that can exceed the 90-day test. The team has one marketer, one SDR and one account executive; existing sales capacity is part of the constraint.

Assume €18,000 media expenditure over the test, 120 distinct valid leads, 36 qualified leads, 12 sales-accepted opportunities and three eventually won customers. Assume all 120 leads belong to one defined acquisition cohort and all later outcomes are observed for that same cohort. The word “eventually” matters: if only 90 days have passed, the wins may still be unknown. Report open opportunities separately and never substitute their full value for realised revenue.

Measure	Formula	Fictional result	Interpretation
Media CPL	€18,000 / 120	€150	Contact acquisition only
Qualified-lead rate	36 / 120	30%	Qualification under the written rule
Cost per opportunity	€18,000 / 12	€1,500	Media cost, excluding sales work
Media CAC	€18,000 / 3	€6,000	Known only after outcomes mature

The assumed opportunity value is $12 \times €12,000 = €144,000$ first-year recurring contract value. That is pipeline scope, not revenue. Three wins imply €36,000 contracted annual recurring revenue before any cancellations or concessions. A finance-approved definition is needed before comparing this value with recognised revenue or cash receipts. The €6,000 media CAC exceeds the €5,000 allowance in the previous model, so the base case would not justify scaling unchanged.

The specific intervention for the pipeline operating system is replace the generic demo promise with a specific workflow assessment and document whether accepted conversations improve. Change that element first, keep the rest of the plan documented, and compare equivalent cohorts after sufficient time. If the next cohort contains only two opportunities, treat its rate as highly uncertain. Write down what additional observation would change the decision, rather than claiming an improvement from one favourable record.

11. Diagnose failure by locating the broken transition

A disappointing aggregate result is a symptom. Diagnose the earliest transition where relevant people fail to progress, then check whether the failure is real or an instrumentation artefact. A low form rate can reflect poor audience fit, an unclear offer, a technical error or a deliberately selective qualification question. Each cause requires a different response. Optimising the visible number without identifying the cause can damage downstream quality.

Begin with record inspection. Select a manageable sample from accepted, rejected and unresolved cases; include records from different campaigns and weeks. Read the actual inquiry and sales notes. Check whether rejection reasons describe a business fact, such as an unsupported use case, or a process fact, such as no one attempting contact. Do not hide a process problem inside a “bad lead” label. Keep an “unknown” category when evidence is missing.

Symptom	First check	Candidate response	Evidence to retain
Many leads, few accepted	Fit and requested next step	Tighten promise or routing	Rejection reason and example
Accepted, no opportunity	Conversation and timing evidence	Improve discovery or decision support	Recorded next step
Pipeline, few wins	Loss reasons and stalled age	Address product or buying barrier	Buyer-confirmed explanation
Platform/CRM gap	IDs, windows and duplicates	Repair measurement	Before/after reconciliation

For the pipeline operating system, inspect whether sales rejection concerns genuine fit, an unclear promise or simply missing follow-up. Do not make three simultaneous creative, targeting and routing changes if you want to learn which cause mattered. Create a diagnosis card: observation; competing explanations; smallest reversible test; owner; cost limit; date of review; decision if the result is inconclusive. A null result is useful when it rules out an expensive assumption.

Stop the affected flow immediately when submissions are lost, exclusions fail or commercial promises are materially misleading. Pause budget expansion when accepted records cannot be followed up within the agreed capacity. Rebuild the offer if the intended segment repeatedly reports that the problem is irrelevant. A healthy operating system makes stopping possible without treating a stopped experiment as a personal failure.

12. Make ownership and handovers operational

Name a single accountable owner for each transition, even when several people contribute. Marketing owns the promise and acquisition context; operations owns dependable movement of records; sales owns acceptance and opportunity evidence; finance owns the commercial reporting definition. Assign a substitute for absences. A shared inbox with no decision owner is not an ownership model.

The handover package should contain the contact's requested action, company and role, source context, relevant use case, qualification evidence, permission state, previous interactions and the next action requested. Separate facts from interpretation. "Asked about a security review on Tuesday" is a fact; "high intent" is an interpretation that needs evidence. Avoid copying unnecessary personal data or speculative profiling into notes.

Agree a service window according to your staffed hours and offer. A demo request might need a different response window from a workbook download. State what happens outside working hours and when sales cannot accept more records. The objective is a credible response, not an arbitrary universal speed target. Let recipients know what to expect when they submit the request.

Transition	Sender must provide	Receiver must record	Escalation
Acquisition to operations	Source and requested action	Successful record creation	Marketing owner
Operations to sales	Fit, context and priority reason	Accept, reject or clarify	Sales lead
Sales to nurture	Reason and permitted next step	New journey or suppression	Journey owner
Closed outcome to reporting	Amount, dates and loss/win reason	Reconciliation status	Finance / RevOps

For the pipeline operating system, the sensitive handover is the point at which marketing context becomes a sales-owned next step. Run a tabletop exercise: one relevant record, one unqualified record, one existing customer and one contact who requests no further communication. Ask each owner to show the next system action. Repair ambiguity before enabling automation.

Use rejection categories that can improve a decision: unsupported segment; no relevant problem; wrong timing with stated review date; duplicate; unreachable after the agreed attempt policy; insufficient evidence. Do not return a record to an endless nurture loop merely because an SDR could not reach the person. Review exceptions weekly and change the process when capacity or the offer changes.

13. Run a 30/60/90-day learning cycle

Treat ninety days as a bounded operating period, not a promise that a long B2B buying process will finish. The first month establishes trustworthy evidence. The second month tests a specific commercial hypothesis. The third month decides which element deserves further investment and which needs redesign. Keep a fixed review date and an explicit cash limit while allowing outcome observation to continue after media activity ends.

Days 1–30: agree the target segment and exclusions; document stage definitions; choose the offer; establish economic assumptions; implement the record chain; and test routing with internal cases. Launch only the smallest coherent programme that can answer the main question. A broad launch across several markets, use cases and channels often consumes the budget without producing enough evidence in any one segment.

Days 31–60: inspect real submissions and sales conversations. Reconcile data, classify rejection reasons and change one high-priority bottleneck. Keep a change log with dates so later cohorts can be interpreted correctly. Interview people who did not progress when that is appropriate and permission permits. Their explanation can reveal a mismatch that no click report will show.

Days 61–90: compare the matured parts of each cohort, examine unresolved opportunities and produce a decision memo. Choose continue, modify, pause or stop. If you continue, state which uncertainty remains and what observation period is needed. If you scale, specify the next budget increment and the capacity that must be available. Do not release the entire annual budget because one month looks promising.

For the pipeline operating system, the first-month deliverable is one agreed segment, one buyer-decision map and a tested lead-to-opportunity trail. Its acceptance test is whether another owner can reproduce the decision from the recorded evidence. Keep a backlog of attractive ideas, but allow only the ideas that answer the current hypothesis into the active test.

Copyable review agenda: what changed; what evidence matured; what remains unknown; where money and staff time went; which failure mode appeared; what will happen next; who owns it; when the decision will be revisited. End with a written decision and a named owner, not a request for more dashboards.

14. Use the workbook and make an evidence-based decision

Use this chapter as a working document rather than a reading exercise. Complete the blanks with actual company evidence, mark unknowns honestly and attach examples where a definition could be disputed. The minimum useful output is one page that lets marketing, sales and finance make the same decision from the same inputs. A larger spreadsheet is justified only when it resolves a real disagreement.

Copyable programme brief

- Business problem: ____; target company and relevant role: ____; explicit exclusions: ____.
- Buyer decision this programme helps: ____; evidence needed to move forward: ____.
- Main hypothesis: if we change ____ for ____, then ____ should improve because ____.
- Acquisition cohort: ____; observation window: ____; mature outcome date: ____.
- Cash limit: ____; staff capacity: ____; economic ceiling: ____; downside assumption: ____.
- Primary decision metric: ____; quality guardrail: ____; data-quality guardrail: ____.
- Stop condition: ____; correction owner: ____; next review date: ____.

Copyable record-review sheet

For each reviewed record capture a pseudonymous reference, acquisition date, source, requested action, company-fit evidence, actual qualification evidence, owner, response outcome, linked opportunity and exclusion reason. Keep sensitive data in the approved CRM rather than exporting a fresh uncontrolled list. Summarise recurring patterns, then link the authorised reviewer back to the system of record.

Decision rules and limitations

Continue when the evidence is relevant and affordable but still immature. Modify when a specific bottleneck has a plausible repair within the remaining learning budget. Pause when measurement or follow-up capacity fails. Stop when the central commercial assumption is contradicted and no bounded corrective test is justified. These are judgement rules, not a statistical guarantee. Document disagreement and the information that would resolve it.

The final question for the pipeline operating system is: Which constraint should receive the next unit of budget or staff time, and what evidence supports that choice?. Answer it with recorded facts and explicitly labelled assumptions. Keep lead volume, qualification, opportunity scope and revenue on separate lines. If the team cannot explain the denominator or the age of an outcome, do not use the number to expand spend.

Primary platform references are provided below and were checked on 24 October 2026. Product interfaces, permissions and plan availability can change. The operating framework in this guide is an original practical synthesis; it is not an official platform methodology or a claim that the vendor endorses these decision rules.

Primary sources and review date

- [Use HubSpot contact and company lifecycle stages](#) — 2026-10-24
- [Understand HubSpot lead scoring](#) — 2026-10-24
- [Campaign URL parameters in GA4](#) — 2026-10-24
- [GA4 Measurement Protocol events reference](#) — 2026-10-24
- [About enhanced conversions for leads](#) — 2026-10-24

- [About qualified leads and converted leads](#) — 2026-10-24
- [Get started with LinkedIn Conversion Tracking](#) — 2026-10-24
- [Manage HubSpot workflow settings](#) — 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 documented target segment with exclusions
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. [Use HubSpot contact and company lifecycle stages](#)
2. [Understand HubSpot lead scoring](#)
3. [Campaign URL parameters in GA4](#)
4. [GA4 Measurement Protocol events reference](#)
5. [About enhanced conversions for leads](#)
6. [About qualified leads and converted leads](#)
7. [Get started with LinkedIn Conversion Tracking](#)
8. [Manage HubSpot workflow settings](#)

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