S.LAB RESEARCH: 2026

Modern Organic Growth Playbook

How to drive high-intent traffic and create durable demand in the AI era

Table of Contents

An evidence-based guide to why B2B traffic is declining and how leading companies build durable demand in the AI era.

Each chapter is designed to be read in sequence, but you can also jump to the sections most relevant to your challenges.

PART 1 — Understanding the Shift

01 What Happened to B2B Traffic 01

02 Assumptions That Didn't Survive the Data 07

03 What Your Traffic Number Hides 15

04 Which Force Hit You 23

PART 2 — Building Durable Demand

05 The Foundation of Durable Demand 33

06 When Growth Starts to Crack 42

07 How Companies Actually Recovered 49

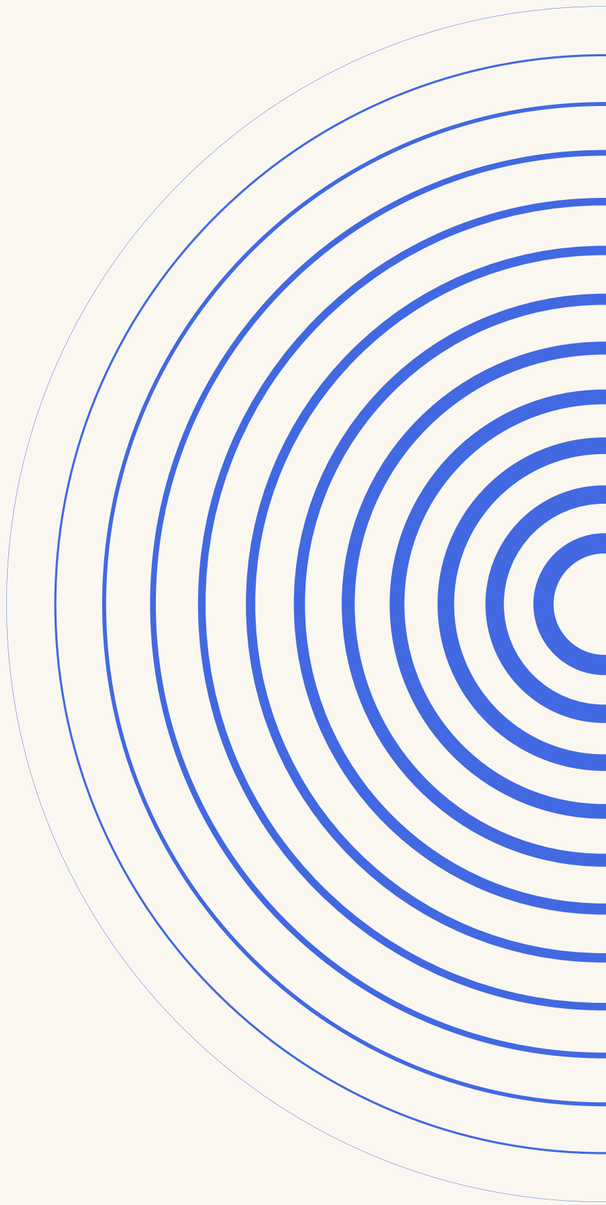
08 When Recovery Isn't the Right Strategy 56

PART 3 — Competing in the AI Era

09 Understanding AI Visibility 65

10 Building for AI Visibility 73

11 The Durable Demand Playbook 79



CHAPTER

01

What Happened to B2B Traffic

B2B organic traffic split decisively in 2025 and 2026 across companies in the same categories, using the same strategies, targeting the same types of queries. Some programs grew by double digits. Some lost half their traffic. Some lost more.

The split wasn't random. It did not map cleanly to company size, publishing volume, domain authority, or obvious differences in effort.

This study tracked 118 B2B SaaS companies across an 18-month window, examining their organic traffic, position data, traffic value, and content composition to discover why programs with similar resources ended up in such different places.

The strongest recurring differentiator wasn't effort, scale, or technical sophistication. It was the type of content each program built.

IN THIS REPORT

01 — What Happened to B2B Traffic

02 — Assumptions That Didn't Survive the Data

03 — What Your Traffic Number Hides

04 — Which Force Hit You

05 — The Foundation of Durable Demand

06 — When Growth Starts to Crack

07 — How Companies Actually Recovered

08 — When Recovery Isn't the Right Strategy

09 — Understanding AI Visibility

10 — Building for AI Visibility

11 — The Durable Demand Playbook

The split

Of the 118 companies tracked, 44 grew organic traffic over the study window, 19 stayed roughly flat, and 55 declined. That distribution — 37% growing, 16% flat, 47% declining — understates how far apart the outcomes were. Many of the winners grew substantially. Many of the losers gave back a meaningful share of traffic they had built over years.

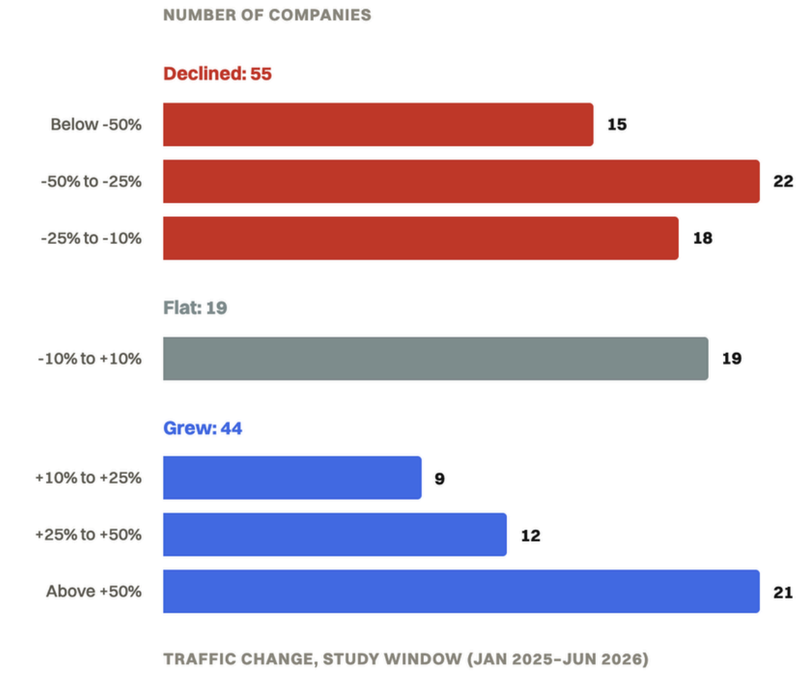
What stands out even more is that the split did not map neatly to industry, company stage, or budget.

Companies in the same category, competing for the same queries, still landed on opposite sides of it. The difference was not where they competed. It was what they had built to compete there.

FIGURE 1.1

The Cohort Split

Traffic change across 118 B2B SaaS companies, January 2025 to June 2026. The cohort divided sharply. 55 declined, 44 grew, and 19 stayed roughly flat. Many fell well beyond modest movement in either direction.



Multiple forces changed the game

Three simultaneous forces produced the split. Understanding them as separate rather than a single “AI disruption” is what makes diagnosis and response possible.

ALGORITHM UPDATES

The August 2025 spam update produced the most identifiable signature: sudden simultaneous collapse across positions 11–20, 21–50, and 51+. Companies anchored in top positions on high-authority pages held or recovered.

AI OVERVIEW INTERCEPTION

At the SERP level, AI Overviews began intercepting queries that could be fully satisfied without a click. Rankings held or improved while traffic fell 80-99% on the same pages.

AI TOOL DEMAND SUBSTITUTION

Furthest upstream, search demand for entire query types declined. Not because AI Overviews answered them in SERP, but because users were generating answers from AI tools before ever searching.



KEY INSIGHT

These forces occurred at the same time but were not the same thing. Each left a different data signature and required a different diagnosis and response. Programs that treated everything as one “AI problem” often applied the wrong fix.



NOTE

These three forces explain most of the broad pressure across the cohort, but they weren’t the only causes of decline. Chapter 4 expands the framework to five decline mechanisms by adding two company-specific patterns that also surfaced repeatedly in the data: architectural damage and competitive displacement.

Content type was the most consistent predictor

Of the companies that grew, nearly two-thirds did so through what we call “compounding content”—engagement-requiring content that accumulates authority and audience over time instead of decaying after an initial traffic spike.

Programs that grew through compounding content had one thing in common: their core traffic assets required continued engagement to deliver value. That ongoing engagement is what allowed demand to continually renew itself over time.

Programs that declined concentrated their traffic in content that provided a direct answer, template, or step-by-step process that a user could extract without engaging further.

This is the difference between **durable demand** and **fragile traffic**.



KEY INSIGHT

Durable demand is built on engagement-requiring content. Because existing readers return to it and new buyers continually need it, the content compounds over time instead of fading after an initial traffic spike.

Fragile traffic is built on content that satisfies. It provides a one-time answer and loses the user the moment they have it.

Examples from the cohort

BRAZE

Braze grew consistently throughout the study window from a wide range of educational content serving practitioners with genuine ongoing learning needs. That content required continued engagement to deliver value, making it resistant to the same environment that damaged other programs.

AHA!

Aha! had one of the most mature content programs in the cohort — extensive investment, genuine category authority, well-executed pages. It still declined substantially. Its investment was concentrated in template and tactical content: the exact types most exposed to AI Overview interception and AI tool demand substitution.

RAMP

Ramp grew by every headline measure, but the headline obscured real fragility underneath. The glossary and free tools drove genuine growth. But Ramp also carried mid-tier and long-tail rankings that the August 2025 spam update damaged — fragility the headline number didn't show.

KEY INSIGHT

Similar traffic trajectories often reflected fundamentally different realities. Looking at growth alone no longer explained why companies succeeded, stalled, or declined.

Before you move on

Before applying any diagnostic framework, spend 30 minutes identifying which of the three forces is most active in your category and program. Each force requires a different response, and treating them as a single problem produces a single response that fits none of them well.

FORCE	CHECK YOUR DATA FOR
Algorithm update damage	Sudden, broad drop in mid-tier positions (11–50) aligned to a known update date
AI Overviews interception	Rankings stable or improving while clicks fall sharply on the same pages
AI tool demand substitution	Search volume declining on affected queries — not just CTR or ranking position



THE BOTTOM LINE

The question isn't whether these forces affected your program. For most B2B content programs in this study window, they did. The question is which force drove your decline, whether the affected traffic was structurally recoverable, and what the right diagnostic sequence is before you spend anything on recovery.

CHAPTER

02

Assumptions That Didn’t Survive the Data

The conventional wisdom about B2B organic growth in 2025 was confident, broadly shared, and frequently wrong.

Several assumptions behind B2B content investment did not survive contact with the data: what AI Overviews suppress, what traffic growth means, where growth comes from, and what actually compounds.

This chapter documents the four assumptions the cohort evidence most directly overturned. Each points to a specific investment error the conventional advice would have produced.

IN THIS REPORT

01 — What Happened to B2B Traffic

02 — Assumptions That Didn’t Survive the Data

03 — What Your Traffic Number Hides

04 — Which Force Hit You

05 — The Foundation of Durable Demand

06 — When Growth Starts to Crack

07 — How Companies Actually Recovered

08 — When Recovery Isn’t the Right Strategy

09 — Understanding AI Visibility

10 — Building for AI Visibility

11 — The Durable Demand Playbook

Where conventional wisdom broke down

FIGURE 2.1

Assumptions vs. Reality

Four beliefs the industry treated as settled going into this study. Each one broke against the data.

WHAT THE INDUSTRY ASSUMED	WHAT THE DATA SHOWED
ASSUMED ~~AI Overviews suppress informational content.~~	DATA SHOWED **Suppression was specific to satisfiable queries. Definitional content grew in five independent categories.**
ASSUMED ~~Traffic growth reflects program health.~~	DATA SHOWED **Traffic and business health decoupled in multiple documented cases. Traffic fell 66% in one case while ARR grew 22%. In another, 154.7% traffic growth corresponded with only 14.7% value growth.**
ASSUMED ~~Concentrated bets drive growth.~~	DATA SHOWED **No growing company showed a page exceeding ~20% of gross gains. Distributed breadth outperformed concentration in every clean-growth case.**
ASSUMED ~~Volume alone compounds traffic growth.~~	DATA SHOWED **Cannibalization rates as high as 286.6% show that volume without portfolio management produces a replacement treadmill, not compounding.**

The most obvious assumptions are often the ones most worth examining. In this study, the beliefs held most confidently by industry were the ones most consistently contradicted by the evidence.

WHAT CHANGED

These findings forced revisions to our working model before the final analysis. The framework used in the rest of this report reflects reality rather than assumptions.

AI Overviews suppress informational content

WHAT THE INDUSTRY BELIEVED

AI Overviews would systematically reduce traffic to informational and educational content: the "what is X" and "how does X work" pages most B2B programs depend on. The theory: if AI can summarize any informational query, the entire category would lose clicks.

WHAT THE DATA SHOWED

Definitional content grew through the AI Overview expansion window across five different categories. AI Overviews did not systematically suppress foundational educational content.

Suppression followed a different logic when it did occur. Content AIO could fully satisfy lost clicks: role lists, templates, step-by-step processes. Content that required engagement to be useful remained largely intact.

THE INVESTMENT ERROR

Retreating from informational content entirely. The programs that pulled back abandoned a strong position based on an assumption the cohort data doesn't support.

CHAPTER 2 — The Second Assumption

Traffic growth always reflects program health



WHAT THE INDUSTRY BELIEVED

A growing organic traffic number means the program is working. More is better.



WHAT THE DATA SHOWED

ZoomInfo's traffic grew 154.7% while traffic value, i.e., what that traffic would cost to acquire through paid search, grew only 14.7%. The headline was among the strongest in the cohort. The underlying composition had shifted toward low-intent, low-value queries.

The misread runs in both directions.

LogRocket's organic traffic sat 66% below its all-time peak while annual recurring revenue grew 22% year-over-year. The business was healthier than traffic suggested. The content program was weaker than revenue suggested.



THE INVESTMENT ERROR

Trusting the headline traffic number. Volume and quality can move in opposite directions. So can traffic and business health. Neither divergence surfaces until you decompose what's actually driving the number.

Concentrated bets drive growth

WHAT THE INDUSTRY BELIEVED

Organic growth comes from breakout assets: a single page, a cluster, or a campaign that achieves outsized rankings on high-volume queries. The strategic goal is to find and scale the hero asset.

WHAT THE DATA SHOWED

Most companies in the study showed the same pattern. No single page exceeded approximately 20% of gross gains. Dozens of pages contributed modest individual gains that added up to substantial aggregate growth.

Programs concentrated in a hero asset faced a different risk profile. Concentration turns a single ranking change, algorithm update, or AI Overview interception into a program-level event. Distribution absorbs those losses across the portfolio.

THE INVESTMENT ERROR

Betting on one page to carry the program. Growth in the strongest programs came from dozens of contributors, not one. The programs that held up didn't look impressive page by page. They looked impressive in total.

Publishing more compounds growth



WHAT THE INDUSTRY BELIEVED

Content is a compounding asset. More pages create more entry points, more rankings, more traffic.



WHAT THE DATA SHOWED

At scale, volume produced diminishing returns. In several cases, it even produced negative returns.

LogRocket's cannibalization rate was 75.6%: for every unit of traffic gained, 0.76 units were lost elsewhere. Aha!'s was 286.6%: for every unit gained, 2.86 were lost.

Both had been producing at significant volume. That production increasingly replaced existing losses rather than building new traffic.

At high cannibalization rates, publishing more doesn't compound the program. It accelerates the treadmill.



THE INVESTMENT ERROR

Publishing more without properly measuring output. At high cannibalization rates, new content replaces lost traffic rather than building it. The program spends to remain stagnant.

Scale doesn’t always protect

The four assumptions explain much of the cohort, but one additional misconception surfaced repeatedly: that scale alone provides protection.

GitHub — the authority case

GitHub’s decline is the clearest counterexample to the assumption that domain authority and scale provide structural protection. GitHub is among the most authoritative domains in developer search, with decades of link equity, massive breadth, and institutional trust. It still declined substantially.

Scale and authority create advantages; they don't create immunity. GitHub likely faced compound mechanisms (AI Overview interception on satisfiable developer queries, demand substitution as AI coding assistants matured) that historical authority couldn't offset. The lesson: authority is necessary but not sufficient when the mechanism driving decline is structural.



WHAT THE DATA SHOWED

Four assumptions commonly used to justify B2B content investment decisions in 2025 were challenged by cohort evidence. In each case, the evidence pointed in a different direction — not universally, but consistently enough to treat each assumption as a prior to examine rather than a foundation to build on.

Before you move on

Before advancing to the diagnostic chapters, run your current content roadmap against each of the four overturned assumptions. A roadmap built on any of these priors will systematically produce the wrong investment decisions.

ASSUMPTION	WHAT TO CHECK
AI broadly suppresses informational content	Check whether the SERP can satisfy those queries without a click. Suppression is query-specific, not category-wide.
Traffic growth = program health	Is the headline traffic number driving decisions? Run the three decompositions in Chapter 3 before acting on it.
Concentrated bets drive growth	Calculate what share of gross gains your top page accounts for. Above 20% is a concentration risk.
Volume compounds	Do you know your cannibalization rate? If not, Chapter 6 walks through the calculation. Run it before expanding production.



THE BOTTOM LINE

The assumptions that felt most obvious were the ones most likely to produce the wrong investment decision. Challenge your original thinking before building an unsustainable content engine.

CHAPTER

03

What Your Traffic Number Hides

The organic traffic number is the most commonly reported metric in content marketing and one of the easiest metrics to misinterpret.

That's because it collapses multiple forces into a single up-or-down figure and strips away most of the context you need.

Across the cohort, headline traffic often gave the wrong read on what was really happening. Growing companies sometimes looked healthier than they were. Declining companies sometimes looked worse than they were. In several cases, the headline number pointed in almost the opposite direction of the underlying data.

Before drawing conclusions, you first have to decompose the number.

IN THIS REPORT

01 — What Happened to B2B Traffic

02 — Assumptions That Didn't Survive the Data

03 — What Your Traffic Number Hides

04 — Which Force Hit You

05 — The Foundation of Durable Demand

06 — When Growth Starts to Crack

07 — How Companies Actually Recovered

08 — When Recovery Isn't the Right Strategy

09 — Understanding AI Visibility

10 — Building for AI Visibility

11 — The Durable Demand Playbook

Why the headline isn’t enough

The headline traffic number answers one question: did total traffic go up or down?

It doesn't answer whether the program is healthy, whether the growth is durable, or whether the decline reflects something the team can address.

A program that grew 40% via brand search and a program that grew 40% via compounding, non-branded content have the same headline and fundamentally different futures.

Three decompositions consistently changed the interpretation of the headline in this study. Each reveals a different layer of what the number hides.



REMEMBER

The headline tells you the score. It doesn't tell you how the points were scored, whether the scoring was sustainable, or whether your opponent was playing the same game. Every strategic decision made from the headline alone carries that uncertainty forward.

Branded vs. non-branded content

Brand traffic grows when more people search for your product by name. Non-branded traffic grows when your content earns rankings on category and educational queries: queries people run before they know your product exists.

Both show up in the headline as organic traffic growth. They are not the same thing.

Brand traffic reflects product adoption; product momentum drives it, not content investment.

Non-branded traffic reflects the content program's ability to capture demand that doesn't know it's looking for you yet. That's the compounding signal.

When brand and non-branded traffic move in different directions, the headline composite can obscure both movements simultaneously.



RUN THIS

Separate your traffic into branded and non-branded queries. If branded accounts for more than 40% of your traffic growth in any given period, the headline is primarily reflecting product momentum, not content program health. They're different levers and require different responses.

Traffic value vs. raw traffic

Traffic value estimates how much it would cost to replace your organic traffic with paid search ads. It weighs traffic by the commercial intent and search volume of the queries driving it, so high-volume, high-CPC keywords contribute more than lower-intent searches.

When raw traffic and traffic value move together, the mix of keywords driving traffic is generally stable. When they diverge, that mix has changed. A site may be gaining lower-value informational traffic, or it may be losing overall traffic while holding onto its highest-value rankings.

A gap of 20 percentage points or more between the raw traffic trend and the traffic value trend is a signal worth investigating before drawing conclusions from the headline.



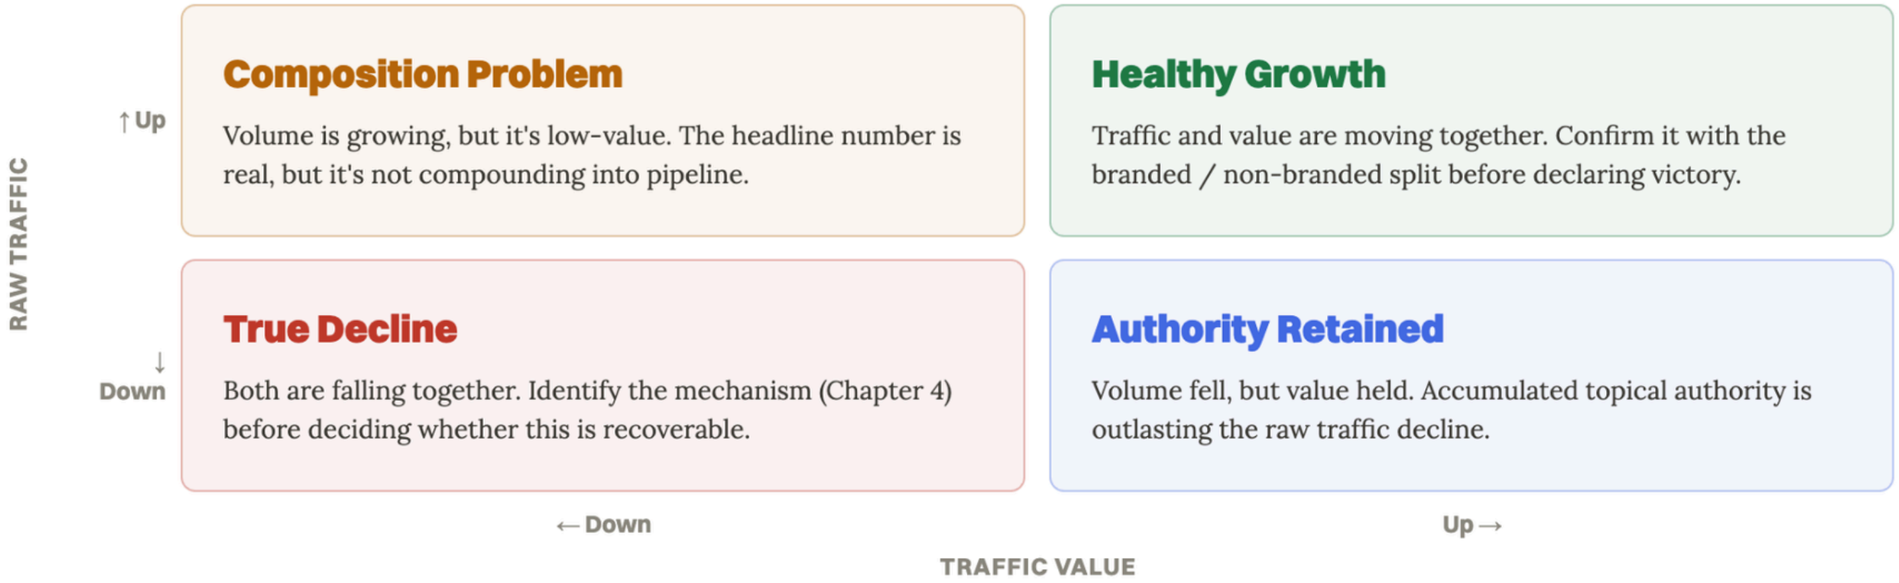
RUN THIS

Pull traffic value and raw traffic from your SEO tool for a 12–18-month window. Calculate the percentage change in each separately, then measure the point difference. A gap of 20 or more points signals a composition shift. Locate your program in the matrix on the next page.

FIGURE 3.1

The Traffic Value Divergence Matrix

Locate your program by crossing raw traffic direction against traffic value direction. Each quadrant implies a different diagnosis.



Category baseline

Organic traffic doesn't move in isolation. It moves within broader category conditions, including algorithm updates, AI Overview expansion, and shifts in search demand.

A 30% decline means something different in a category where the median company declined 25% than in one where the median grew 15%. The absolute number tells you what happened. Comparing your results to the rest of the category tells you whether your performance was typical or unusual.

Context matters in both directions. A decline that looks severe on its own may actually be better than the rest of the category. Likewise, strong growth may simply reflect a category-wide upswing. Without that comparison, it's easy to misdiagnose the problem.

Category-wide changes are often hard to recognize from inside a single program. Algorithm updates, AI Overview expansion, and changes in search behavior can affect entire categories at once, making external forces look like problems unique to your content.



WHAT TO DO

Identify one or two direct category peers. Pull their traffic trend for the same window and compare directionally. If they declined similarly, the headwind is likely market-wide. If they held or grew while you declined, the problem is specific to your program.

Examples from the cohort

Gainsight (Branded vs. Non-Branded)

Roughly 65% of Gainsight's traffic growth came from a single branded keyword. The headline showed a growing program. Underneath: limited non-branded content, carried by navigational searches from an existing customer base. The headline and the content program were telling different stories.

ZoomInfo (Traffic Value vs. Raw Traffic)

ZoomInfo produced the most extreme divergence in the study: traffic grew 154.7% while traffic value grew only 14.7%. Most of the volume growth was in low-value, low-intent queries. The headline implied the program had more than doubled in effectiveness. The traffic value data disagreed.

Freshworks (Category Baseline)

Freshworks' traffic decline looked like a program-level failure in isolation. The picture changed against the category baseline. The CS and CRM category declined sharply during the same period. The headline suggested a failing content program. The category comparison suggested a market-wide headwind.



WHAT THE DATA SHOWED

In most cases studied, the headline number gave an incomplete or misleading picture of the organic program. The three decompositions consistently changed the interpretation of that number. The implication is simple: any strategic decision made from the headline alone is made with missing context.

Before you move on

Run these three checks before making any strategic content decisions:

ASSUMPTION	WHAT TO CHECK
Branded vs. Non-Branded Split	Is the traffic growth actually coming from content, or mostly from brand demand?
Traffic Value vs. Raw Traffic	Is traffic quality moving with traffic volume, or are they diverging?
Category Baseline	Is this change specific to our program, or is the whole category moving the same way?

Once you have a verified picture of what your traffic number actually measures, the next question is: which specific mechanism caused what you see?



THE BOTTOM LINE

The headline traffic number tells you whether the composite went up or down. It doesn't tell you whether the program is healthy, whether the growth is durable, or whether you're diagnosing the right problem.

CHAPTER

04

Which Force Hit You

Most content teams respond to traffic decline the same way: publish more content, build more links, refresh existing pages.

There's just one problem: those are the right responses to only one of the five mechanisms this study identified.

If you misdiagnose the mechanism, even good work becomes wasted work.

IN THIS REPORT

01 — What Happened to B2B Traffic

02 — Assumptions That Didn't Survive the Data

03 — What Your Traffic Number Hides

04 — Which Force Hit You

05 — The Foundation of Durable Demand

06 — When Growth Starts to Crack

07 — How Companies Actually Recovered

08 — When Recovery Isn't the Right Strategy

09 — Understanding AI Visibility

10 — Building for AI Visibility

11 — The Durable Demand Playbook

Start with your trajectory

Across the 118 companies in this study, traffic moved in recognizable patterns. Five distinct trajectory types appeared consistently:

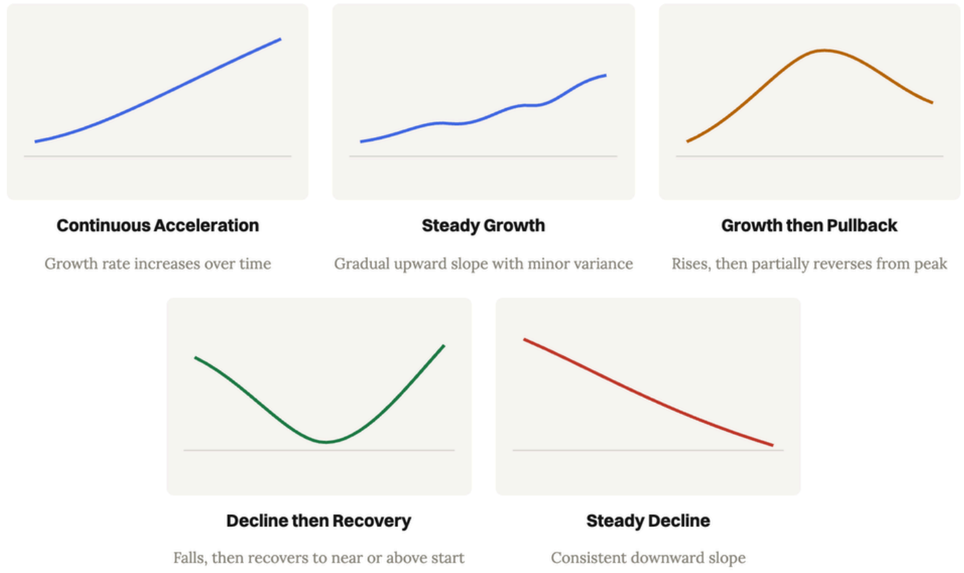
1. Continuous Acceleration
2. Steady Growth
3. Growth then Pullback
4. Decline then Recovery
5. Steady Decline

Trajectory tells you the shape of what happened. It does not tell you why. Two companies can show the same “Decline then Recovery” curve for different reasons: one because an algorithm update reversed, another because brand demand expanded into new geographies. Position data, timing, and search volume reveal the cause.

FIGURE 4.1

Five Trajectory Types

The shapes traffic took across the study window, before any mechanism is attached to them.



Where do you fall?

Map your most visible symptom to the most likely mechanism. Use the full diagnostic descriptions to confirm.

IF YOU SEE...	SUSPECT...
Rankings stable, traffic collapses by 80-99%	**AI Overviews interception**
Sudden collapse across multiple position buckets, aligned to an update	**Algorithm update damage**
Traffic drop begins the same week as a redesign or migration	**Architectural damage**
Search volume itself declining on affected queries	**AI tool demand substitution**
Slow position erosion over months, query by query	**Competitive displacement**



KEY INSIGHT

Publishing more content is the right response to competitive displacement. It is not the right response to architectural damage, AI Overview interception, algorithm updates, or upstream demand substitution.

The right move depends on the root cause — a diagnosis most programs skip.

AI Overview Interception



SIGNATURE

Position holds or improves while traffic collapses 80–99%.



LOOK FOR

- Stable or improving rankings on the same keywords where traffic fell
- The search result itself answers the affected query: role lists, templates, brief definitions, step-by-step processes
- CTR collapsed at the same position



CASES

Aha! showed 20+ queries where pages held or improved rankings while losing 80–99% of traffic. Chapter 9 covers this in detail.



DON'T CONFUSE WITH

Competitive displacement. That's when position and traffic fall together. Under AIO interception, position holds or improves while traffic disappears. If position is also declining, suspect competitive displacement instead.



PRIMARY RESPONSE

Don't refresh by default. First, ask whether the page gives the user something they can't get from the Overview. If the value can be extracted and summarized, the query is structurally exposed. If the value requires deeper engagement, the page has more protection.

Algorithm update damage



SIGNATURE

Sudden, simultaneous collapse across positions 11-20, 21-50, and 51+, within the same short window, aligned to a known algorithm update date.



LOOK FOR

- Multi-bucket collapse in days, not months
- Timing matched to a published algorithm update
- Top positions (1-10) holding or strengthening while mid-tier collapses



CASES

Ramp and Cloudbees (August 2025 spam update)



DON'T CONFUSE WITH

Competitive displacement, which erodes positions gradually over months, one page or query at a time. If the week everything changed aligns to a known update, algorithm damage is more likely.



PRIMARY RESPONSE

Identify what the update appears to have targeted before investing in recovery. Algorithm losses can be reversible, but refreshing pages before you know what changed usually wastes time.

Architectural damage

SIGNATURE

Sharp, sudden decline closely aligned to a dateable internal event: a redesign, platform migration, CMS change, or content overhaul.

LOOK FOR

- Timing is unusually clean: decline begins within the same week or month as a known site change
- Backlinks and domain authority remain intact, but rankings fall
- No corresponding algorithm update in the same window

CASES

Freshworks: December 2025 "Dew" design system launch. Referring domains were retained; rankings were not.

DON'T CONFUSE WITH

Algorithm update damage, which is also sudden but correlates to external update dates, not internal events. If a specific site change coincides with the timing of decline, and no major algorithm update occurs in the same window, architectural damage is the most likely mechanism.

PRIMARY RESPONSE

Fix the architecture before publishing anything new. The most common culprit: on-page content changes that altered targeting signals even when URLs were preserved and redirects were correctly configured.

AI tool demand substitution

SIGNATURE

Search volume on the affected queries declined. Not just CTR or ranking: the queries themselves are happening less frequently.

LOOK FOR

- Monthly search volume falling over the past 12-18 months
- Affected content types: templates, how-to content for routine tasks, developer tutorials for mature technologies
- Decline accompanied by AI tool adoption in the same category

CASES

Aha! lost roughly 60% of traffic across 154 pages of template content. LogRocket's developer tutorial traffic declined alongside the maturation of AI coding assistants.

DON'T CONFUSE WITH

AIO interception, which operates at the SERP level. Search still occurs, but the click doesn't. Demand substitution happens before the search: the user generates the artifact or answer directly without searching. If volume data shows declining search frequency, the mechanism is upstream of the SERP.

PRIMARY RESPONSE

Stop investing in the affected content type for SEO purposes. Demand that has migrated to AI generation doesn't return through content refresh or link building. Redirect investment toward query types with structurally durable demand.

Competitive displacement

SIGNATURE

Gradual, sustained position erosion over months across specific queries, as competitors outrank you on individual pages.

LOOK FOR

- Decline is measured in months, not weeks
- Position and traffic fall together
- You can identify the specific page that overtook yours and observe when it rose as yours fell

CASES

Competitive displacement rarely produces a distinctive data signature. You identify it after the more specific mechanisms do not fit. The decision tree saves it for last for exactly that reason.

DON'T CONFUSE WITH

Algorithm update damage or AIO interception. Both can produce traffic loss without obvious ranking shifts in simplified dashboards. Confirm by finding the page that outranked yours and observing when it rose as yours fell.

PRIMARY RESPONSE

Publish better content or earn more authoritative links on the affected queries. Competitive displacement is the one mechanism where "produce more content" is actually the correct diagnostic conclusion.

Running the diagnostic sequence

"Sudden" means a visible cliff in the data. Gradual declines go directly to competitive displacement — confirm with position data.

For sudden declines, split by event type: internal (something your team changed) or external (a published algorithm update).

Internal event, no corresponding algorithm update → architectural damage. Check backlink data as a control.

Algorithm update → check position buckets. Multi-bucket collapse across 11–20, 21–50, and 51+ in a short window is the algorithm damage signature.

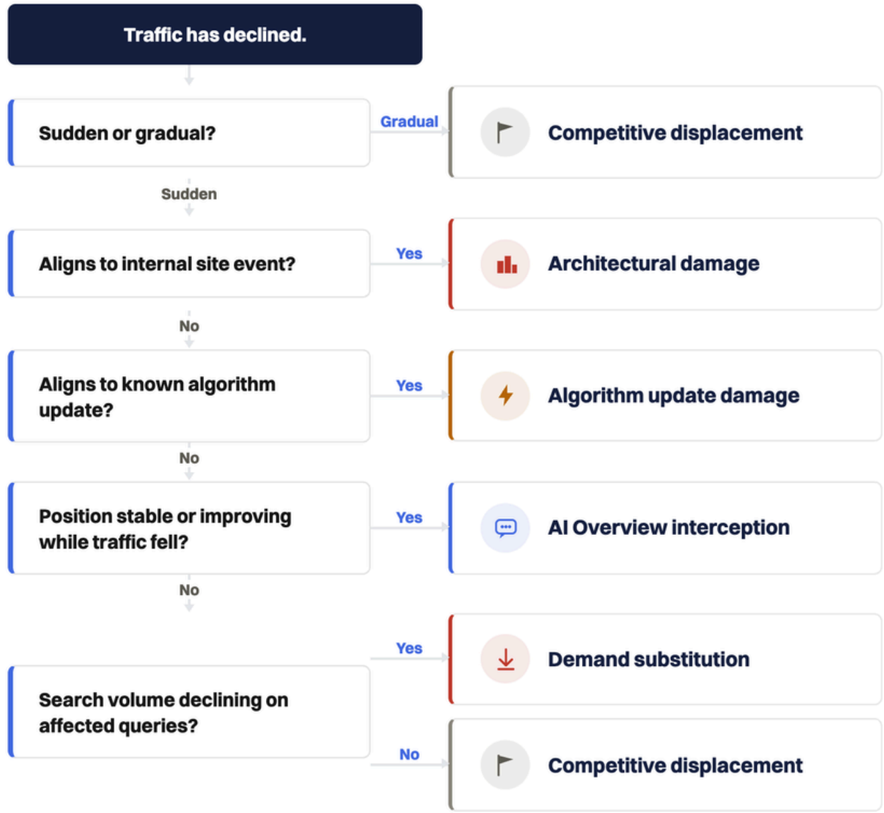
Neither → compare position and traffic. If positions hold or improve while traffic falls, suspect AIO interception. Check whether the affected queries can be answered without a click.

In every case: if search volume on affected queries has declined, demand substitution is operating upstream.

FIGURE 4.2

Diagnostic Decision Tree

Five questions to identify which mechanism is driving a traffic decline before spending anything on recovery.



Before you move on

Before committing to any recovery investment, spend 2–3 hours running the diagnostic sequence.

1. Use the symptom table at the top of this chapter to identify the most plausible mechanism before going deeper.

2. Pull position data for your affected pages across the decline window. Note when the decline began. Cross-reference that date against algorithm updates and any internal site changes. Then check keyword volume trends on the affected queries and run the decision tree.

3. Once you've identified the most likely mechanism, find that mechanism's section in this chapter and read the primary response before doing anything else.

The wrong treatment applied with high confidence is the most expensive outcome.



THE BOTTOM LINE

Traffic data tells you what happened. Position data, timing, and search volume tell you why. Get the diagnosis right before you spend money trying to recover.

Once you identify the mechanism, the next step is fixing it. Chapter 5 examines the structural traits of the programs that kept growing through the same environment and why that pattern isn't always what you might suspect.

CHAPTER

05

The Foundation of Durable Demand

The dominant B2B content strategy playbook assumes growth is something you discover. Somewhere in the search landscape, in some topic, format, or angle, there is a move that will produce outsized results if you identify it and execute it well enough. The strategist's job is to find that move and scale it.

But among content-driven growers in this study, **no single page accounted for more than roughly 20% of gains**. Growth came from a broad portfolio of content, anchored in a type of content many teams had started to dismiss.

To understand what drove that growth, you need to look at three things: how gains were distributed, what kind of content produced them, and what position patterns signal durable versus fragile momentum.

IN THIS REPORT

01 — What Happened to B2B Traffic

02 — Assumptions That Didn't Survive the Data

03 — What Your Traffic Number Hides

04 — Which Force Hit You

05 — The Foundation of Durable Demand

06 — When Growth Starts to Crack

07 — How Companies Actually Recovered

08 — When Recovery Isn't the Right Strategy

09 — Understanding AI Visibility

10 — Building for AI Visibility

11 — The Durable Demand Playbook

Distributed growth outlasts concentrated growth

Most content programs follow the same pattern: one or two pages generate most of the traffic, while everything else contributes much less. Teams often see this as success. A breakout page becomes proof that the content program is working.

No single page dominated among the fastest-growing companies in this study. The biggest winner typically accounted for just 15–18% of total gains. The next several pages each contributed another 8–12%. Beyond that, dozens of pages delivered smaller gains that added up over time.

We saw this pattern across DevOps, finance, travel, security, and marketing automation. Growth came from many solid performers, not a single breakout page.

A breakout page feels like success because it's highly visible. Distributed growth is quieter, but it's more resilient. When one page lost traffic, dozens of others grew, keeping the overall program strong.



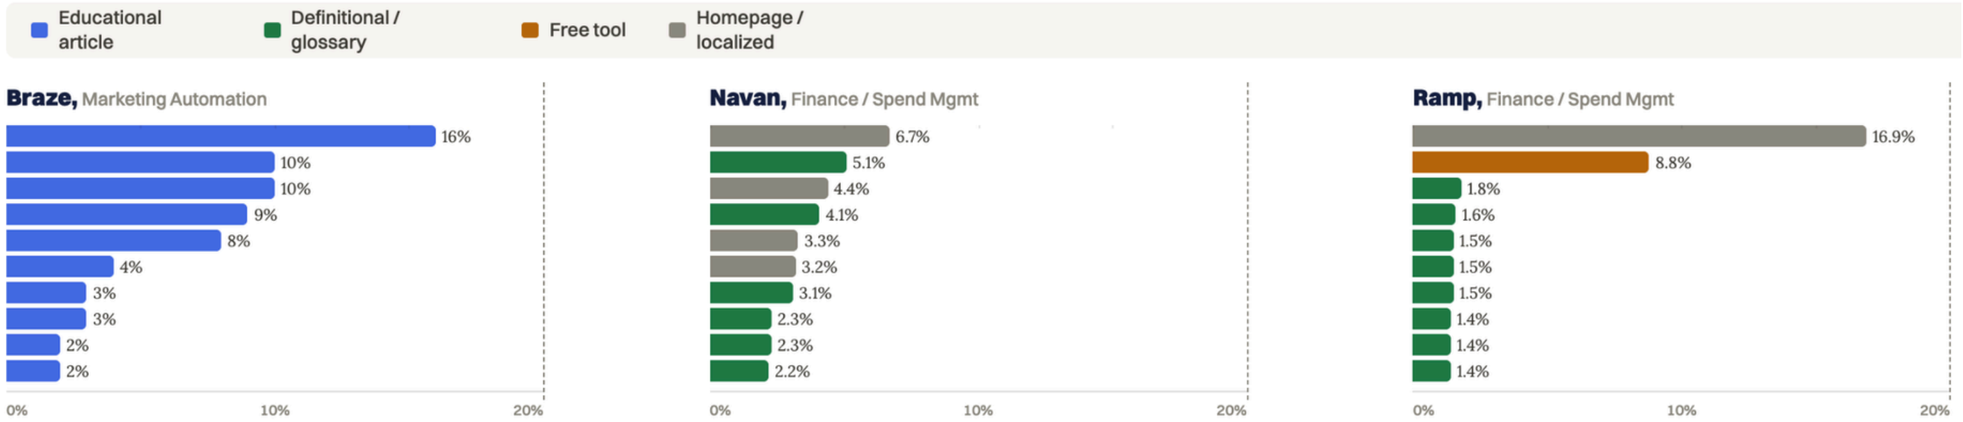
KEY INSIGHT

A program built around one breakout page is a program with one point of failure. When that page loses a ranking, the whole program feels it. Distributed programs absorb losses. Concentrated programs amplify them

FIGURE 5.1

No Single Page Drove More Than ~20% of Growth

Top 10 pages by traffic gain for three content-driven growing companies in the study cohort. Each bar represents one page's share of total gross gains during the study window.



Definitional content captures recurring demand

The distribution exists because of the type of content behind it. Across the winning cohort, the pages generating the most consistent gains were definitional. They answered questions about concepts rather than products, tactics, or workflows:

What is continuous integration? What is SOC 2? What is accounts payable?

These pages help readers understand a concept well enough to leave with a working understanding of it.

Their advantage isn't simply that the concepts don't change very often. It's that the audience keeps renewing itself. Every year, new engineers need to learn continuous integration. Every year, companies prepare for their first SOC 2 audit. The tools evolve, but new people continue searching for the underlying concepts.



KEY INSIGHT

The question isn't whether content is informational. It's whether a reader can accomplish their goal without leaving the AI response. If they can, that content is structurally at risk, regardless of how well it's executed.

Not all informational content is the same

Aha! had a large, mature content program and still declined significantly. The difference was the type of content the program emphasized.

Much of Aha!'s library focused on strategy: roadmapping frameworks, prioritization methods, and product leadership. That's valuable informational content, but readers can often accomplish their goal from an AI-generated summary without reading the full article.

The definitional content that consistently grew across the winning cohort served a different purpose. Understanding what continuous integration or SOC 2 actually means requires more than a summary. Readers need a working understanding of the concept.

That's the distinction captured by the satisfiability test. Two pieces of content may both be informational, but if one satisfies the reader's goal without a click and the other doesn't, their exposure to AI displacement is fundamentally different.



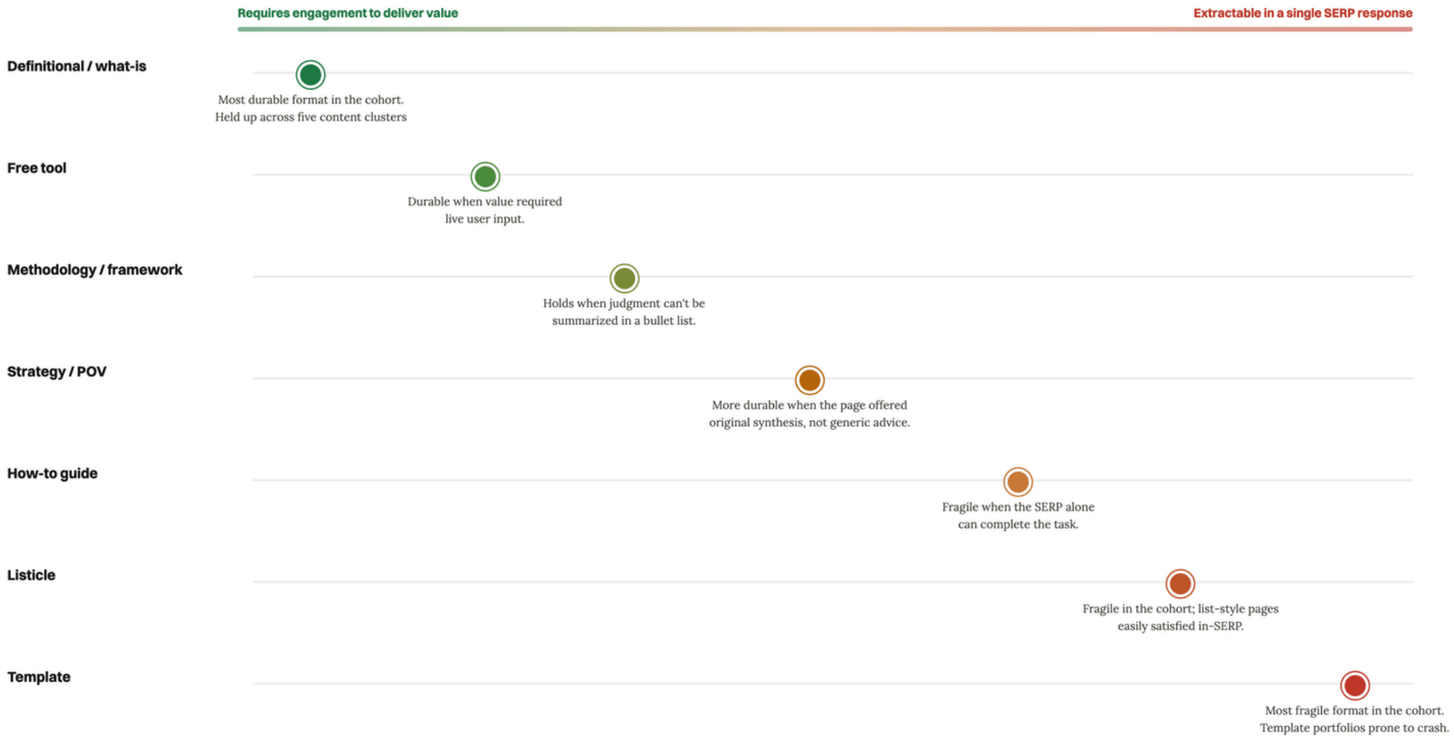
KEY INSIGHT

The durability of definitional content isn't about format or length. It's about whether readers can accomplish their goal without leaving the AI response. If they can, the content is structurally more exposed to AI displacement.

FIGURE 5.2

The Content Durability Spectrum

Seven content types positioned by how they performed in the cohort. What matters is not topic, but whether the content requires engagement to deliver value or can be extracted directly in the SERP.



Across the strongest-performing programs, rankings became increasingly concentrated near the top of the SERP. Rather than maintaining visibility across hundreds of mid-tier keywords, they replaced weaker rankings with stronger positions on the terms that mattered most. Looking only at total keywords obscured this shift.

Strong programs consolidate toward the top of the SERP

Across most growing companies, rankings outside the top 20 declined while top-10 rankings held steady or improved. Rather than losing topical coverage, these companies replaced weak rankings with stronger ones that drove meaningful traffic. They didn't necessarily need *more* keywords as long as they ranked for the *right* keywords.

	OLD METRIC	BETTER METRIC
What to Track	Total ranking keywords	Share of rankings in positions 1-10
Healthy Growth	More keywords ranking anywhere in the SERP	More top-10 rankings, even if total rankings decline
Warning Sign	Keywords dropping out of the index entirely	Keyword count grows, but top 10 rankings don't



KEY INSIGHT

The goal isn't to rank for more keywords. It's to move more keywords into positions that drive meaningful traffic.

Examples from the cohort

Braze (distributed gains)

Braze grew while most marketing automation peers declined. No single page accounted for more than 15–18% of total gains. The next several pages each contributed another 8–12%. Beyond that, dozens of pages delivered smaller gains that added up over time. There was no hero asset. Growth came from a distributed portfolio that didn't depend on any single page.

Vanta (definitional content)

Vanta's growth was anchored in definitional security and compliance content. Pages like "What is SOC 2?" drew consistent traffic because new companies entered the SOC 2 process every year. The demand wasn't driven by news or product launches. It continually renewed itself.

Harness (position consolidation)

During the study window, Harness lost rankings outside the top 10 while its top-10 rankings held or improved. Total keyword count made the program look like it was shrinking. In reality, it was consolidating visibility into the positions that drive meaningful traffic.



WHAT THE DATA SHOWED

Among 38 fully classifiable growing companies, 23 (61%) grew primarily through distributed definitional content. It was the only growth pattern that appeared across all five trajectory types. Brand-led growth and niche concentration appeared only in specific trajectories.

Before you move on

Pull your Top Pages report and calculate what share of gross gains over the past 12 months came from your top-performing page. Among the strongest growers in this study, no single page accounted for more than roughly 20% of gross gains. If one page drives substantially more than that, ask what happens if it loses its ranking.

Then identify two or three topics adjacent to your strongest performers where you have little or no content. Those are your first opportunities to diversify the program.



THE BOTTOM LINE

The question isn't whether you have a breakout page. It's whether your program could absorb losing it. If one or two pages determine your performance, the structure is the problem, not the execution.

CHAPTER

06

When Growth Starts to Crack

Everyone measures net organic traffic growth. Almost nobody measures what had to be lost to achieve it.

That omission is understandable in small programs. In a 50-page content program, you can still track page performance directly. But once a program scales into the hundreds or thousands of pages, a single traffic number hides more than it reveals.

A program that grew 30% in net traffic last year may have generated 200% gross gains while absorbing 170% gross losses. Another may have grown 30% with 35% gross gains and 5% gross losses. Same headline. Completely different health.

Across the growing companies, **three hidden risks appeared consistently**. None is visible in the headline number. All change what that growth actually means.

IN THIS REPORT

01 — What Happened to B2B Traffic

02 — Assumptions That Didn't Survive the Data

03 — What Your Traffic Number Hides

04 — Which Force Hit You

05 — The Foundation of Durable Demand

06 — When Growth Starts to Crack

07 — How Companies Actually Recovered

08 — When Recovery Isn't the Right Strategy

09 — Understanding AI Visibility

10 — Building for AI Visibility

11 — The Durable Demand Playbook

CHAPTER 6 — The First Hidden Risk

Net growth hides gross movement

Portfolio traffic reflects two things happening at once: some pages gain traffic while others lose it. Most reporting combines those movements into a single net number, hiding both.

Four numbers matter:

- **Gross gain:** Total traffic gained across pages that grew.
- **Gross loss:** Total traffic lost across pages that declined.
- **Net gain:** Gross gain minus gross loss.
- **Cannibalization rate:** Gross loss divided by gross gain. In other words, how much traffic the program lost for every unit it gained.

A cannibalization rate of 10% means that for every unit of traffic the program gained, it lost 0.10 elsewhere. At 75%, it lost 0.75. At 286%, it lost 2.86.



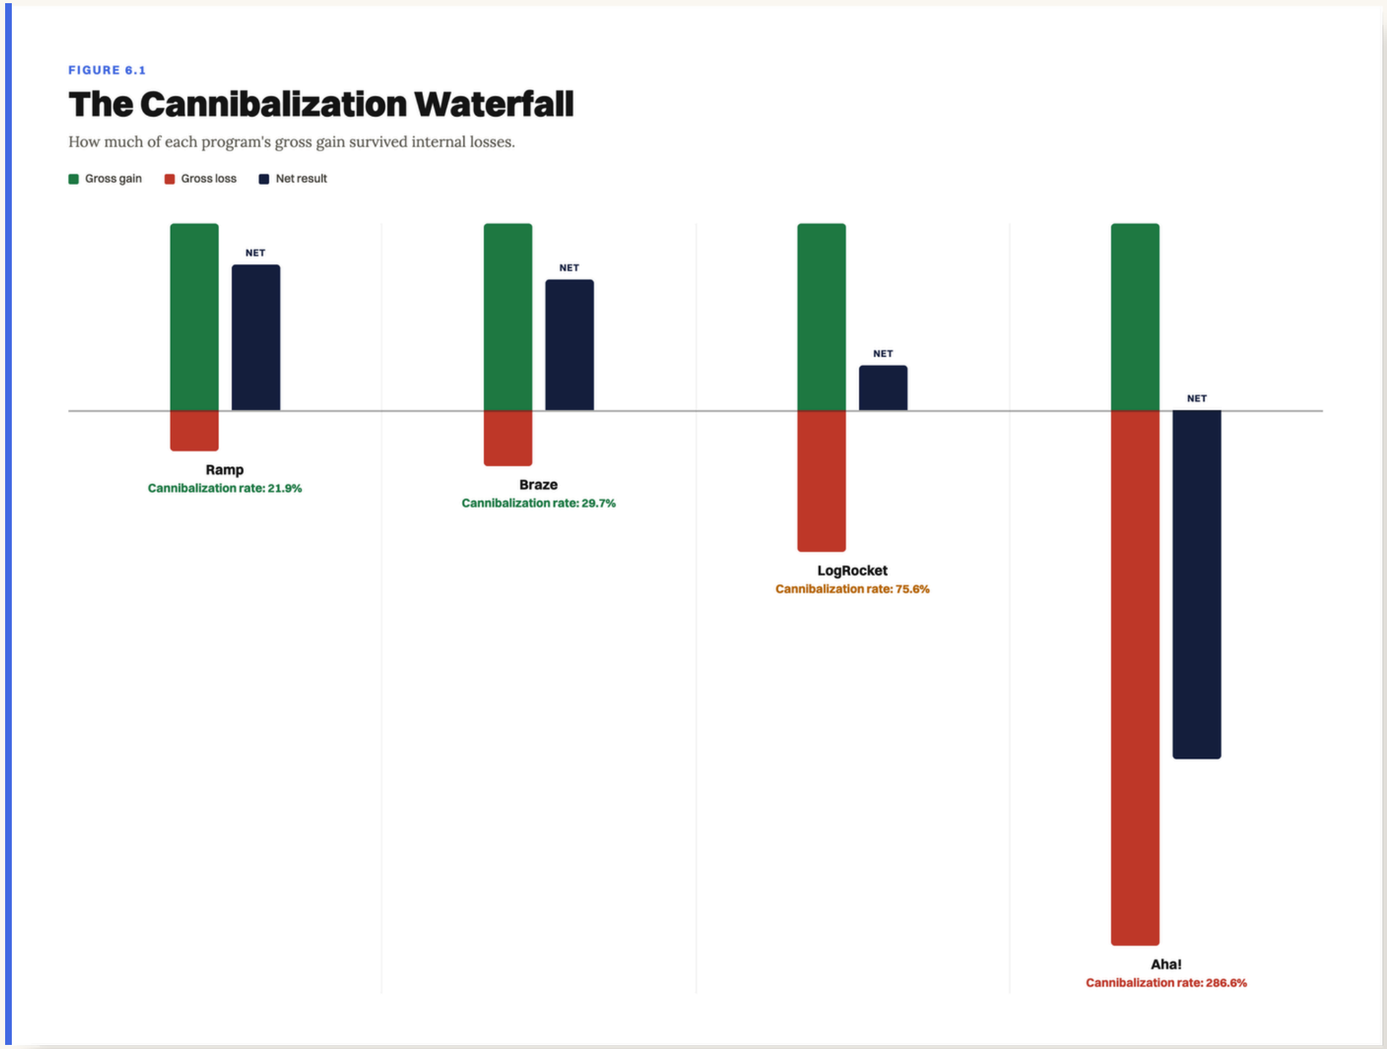
KEY INSIGHT

Net growth tells you the score. Cannibalization rate tells you whether you're winning or just running faster to stay even.

The content treadmill

Cannibalization rate is not just a snapshot metric. It compounds. A program with high cannibalization has to generate far more gross gain to produce a modest amount of net growth, and as the portfolio gets larger, so does the base it can cannibalize.

This is the treadmill. New wins have to replace old losses before the program can grow at all, so more and more effort goes toward maintaining existing traffic rather than building new traffic.



The four cannibalization zones

Healthy: Under 30%

The program is genuinely compounding. Gross gains substantially exceed gross losses, so new content produces durable net growth rather than replacement traffic.

Watch: 30-60%

The program is working, but losses offset a meaningful share of gains. This is worth monitoring, especially if the rate is rising over time.

Treadmill: 60-100%

The program is working much harder than the headline suggests. More than half of every gross gain disappears elsewhere. At 75–80%, producing net growth takes roughly four times the apparent effort.

Net negative: Above 100%

The program is losing more traffic than it creates. The headline may still show growth if gross gains are large enough, but the economics are inverted: each new content win is offset by even larger losses elsewhere.



TO CALCULATE CANNIBALIZATION RATE

1. Export your top pages for two periods.
2. Sum the traffic gains across pages that increased.
3. Sum the traffic losses across pages that declined.
4. Divide the gross loss by the gross gain.

CHAPTER 6 — The Second Hidden Risk

Position collapse: the hidden risk inside headline growth

The second hidden risk appears in the ranking profile rather than the traffic profile.

Ramp grew strongly by headline traffic during the study window. Underneath that growth, however, rankings in positions 21–50 and 51+ collapsed around the August 2025 spam update. Top-performing pages continued gaining traffic, but the broader ranking base that often feeds future top-10 rankings shrank sharply.

At first glance, that can look like the healthy consolidation described in Chapter 5. It wasn't. Healthy programs gradually replace weaker rankings with stronger ones near the top of the SERP. Ramp didn't strengthen its ranking profile. It lost a large share of its long tail after an external shock.

Ramp is one example, not a universal pattern. But it illustrates what headline traffic can hide: a program that continues growing while the foundation for future growth erodes underneath it.



THE DIAGNOSTIC QUESTION

When the long tail shrinks, ask one question: Did stronger rankings replace it, or did it simply disappear? The answer determines whether you're looking at healthy consolidation or structural erosion.

CHAPTER 6 — The Third Hidden Risk

When product momentum hides content performance

The third hidden risk is the easiest to miss because it looks like success.

Brand traffic grows when more people search for your product by name. That's a real organic signal, but it reflects product momentum more than content performance.

Non-branded traffic provides the clearest view of whether the content program is actually capturing demand. When branded traffic grows while non-branded traffic stalls or declines, the headline can look healthy even though the content program isn't driving the result.

The simplest diagnostic is to track branded and non-branded traffic separately over time. If non-branded traffic drives growth, the content program compounds. If branded traffic carries the gains, the headline reflects product momentum rather than content performance, and that momentum may not last.



WHAT TO WATCH FOR

Headline growth only reflects content performance when non-branded traffic drives it. Separate branded and non-branded traffic before judging whether the content program is actually compounding.

Before you move on

Run this three-part audit and compare your results against the benchmarks below.

CHECK	DIAGNOSTIC	HEALTHY	WARNING SIGN
Cannibalization Rate	Gross loss / gross gain (top pages, two periods)	Under 30%	Over 60%
Position Profile	Compare share of keywords in positions 1-20 vs 21+	Top-10 share growing	Sudden long-tail drop after external event
Brand Dependency	Compare branded vs. non-branded traffic	Non-branded drives growth	Branded drives over 40% of growth

*30-60% *cannibalization is a watch zone. Monitor the trend over time.*



THE BOTTOM LINE

Healthy headline growth isn't enough. It must be durable. A few minutes spent checking cannibalization, ranking profile, and brand dependency can reveal risks that total organic traffic will never show.

CHAPTER

07

How Companies Actually Recovered

When traffic declines, the standard prescription writes itself: audit what was lost, identify the gaps, and publish content to fill them. Recovery, in this framing, is a production problem.

But the companies in this study didn't rebuild growth through new content. **Companies recovered because the structural conditions for growth survived the decline.** Existing content did something it hadn't done before: either a cluster consolidated or a single page broke through because the top-position floor held.

Where the losses fell determined whether recovery was possible and which archetype applied.

IN THIS REPORT

01 — What Happened to B2B Traffic

02 — Assumptions That Didn't Survive the Data

03 — What Your Traffic Number Hides

04 — Which Force Hit You

05 — The Foundation of Durable Demand

06 — When Growth Starts to Crack

07 — How Companies Actually Recovered

08 — When Recovery Isn't the Right Strategy

09 — Understanding AI Visibility

10 — Building for AI Visibility

11 — The Durable Demand Playbook

Recovery starts with one question

What survived the decline?

Recovery didn't depend on how much traffic a program lost. It depended on whether the program's structural foundation survived.

Across the recovering companies, losses concentrated in peripheral content, including supporting blog posts, comparison pages, and mid-tier educational content. Meanwhile, foundational assets held, including definitional pages, product-adjacent authority content, and the pages already anchoring top positions.

The non-recovering companies showed the opposite pattern. They lost rankings on pages central to their topical authority. Freshworks lost core customer experience content after its architectural disruption, while Secureframe and Sprinto lost foundational compliance pages such as "What is SOC 2?" and "What is ISO 27001?" These were exactly the kinds of assets that supported durable growth elsewhere in the cohort.

Recovery depended less on the size of the decline than on what survived it. Peripheral losses left the structural floor intact. Foundational losses broke it.



WHAT TO LOOK FOR

Don't start by measuring how much traffic disappeared. Start by identifying which pages disappeared. The location of the loss is often more important than its size.

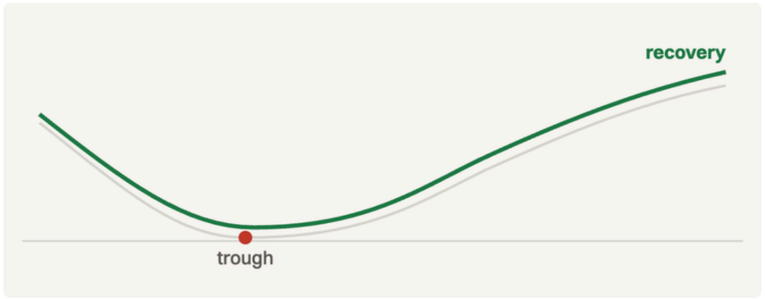
FIGURE 7.1

The Two Recovery Archetypes

Recovering companies pulled recovery from content they already had. What differed was which asset did the lifting.

ARCHETYPE 1

Cluster Breakthrough



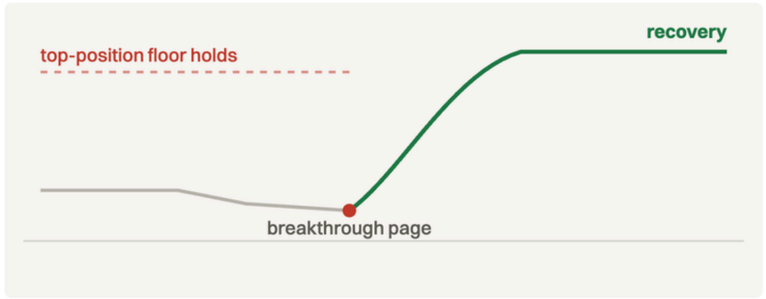
Gradual lift across multiple related pages, rising together.

WHAT HAD TO SURVIVE

Core pages held position through the decline while losses concentrated in peripheral content.

ARCHETYPE 2

Top-Position Floor



Faster recovery driven by one page already in striking range.

WHAT HAD TO SURVIVE

Top-position content (positions 1–3) held while mid-tier and long-tail traffic collapsed.

Don't match either? Go to Chapter 8's structural vs. tactical test before investing in recovery.

Not every recovery came from publishing more content. Some of the strongest recoveries occurred when existing clusters regained visibility together after structural weaknesses were addressed.

Content cluster breakthrough



WHAT SURVIVED

Core foundational pages held their positions throughout the decline. Losses concentrated in peripheral and supporting content, leaving the structural foundation intact.



HOW RECOVERY HAPPENED

Multiple existing definitional and comparison pages improved together. Rankings rose across the cluster without new content production or detectable link acquisition. These pages captured recurring practitioner demand, allowing the entire cluster to strengthen over time.



CASE: CircleCI

CircleCI recovered through a content cluster it had already built, not new investment. Multiple DevOps concept pages rose together, strengthening its presence on foundational educational queries. The recovery was possible because that cluster already existed before the decline.



TYPICAL TIMELINE

6+ months from trough to visible recovery. Recovery emerged gradually as multiple pages improved together.



THE DIAGNOSTIC QUESTION

Do you already have the cluster? Recovery becomes possible when multiple related pages already exist and are close enough to improve together.

Top-position floor

WHAT SURVIVED

Top-position content (positions 1–3) held throughout the mid-tier and long-tail collapse. The program's highest-authority pages never lost their footing.

HOW RECOVERY HAPPENED

Because the top-position floor held, the program didn't need a broad recovery across the portfolio. One existing page broke into a materially higher ranking and drove the recovery.

CASE: Hyperproof

A single ranking breakthrough on "risk management frameworks" (estimated 56k monthly search volume in Ahrefs, position 21 → 8) drove the recovery. The page already existed. It simply moved from out of reach to high visibility without new production.

TYPICAL TIMELINE

2–4 months, the fastest archetype in the study. The floor never broke, so Hyperproof did not need to rebuild from the bottom. It needed one breakthrough on a page already within striking distance.

KEY INSIGHT

When the top-position floor survives, recovery doesn't require the entire program to improve. One breakthrough page can change the trajectory.

Matching the archetype to your situation

Now determine whether either archetype applies to your program. Run this sequence before committing to a recovery plan.

STEP 1: Identify where losses concentrated

Check: Did losses hit peripheral content or foundational assets?

If foundational pages declined, neither recovery archetype applied in this study. When the structural floor breaks, recovery becomes a different problem. Continue to Chapter 8.

If losses were mostly peripheral, continue to Step 2.

STEP 2: Assess the top-position floor

Check: Did positions 1–3 hold during the decline?

If they did, assess Archetype 2. Look for existing pages ranking in positions 8–25 on high-value queries.

If top positions also fell, move to Step 3.

STEP 3: Look for existing recovery opportunities

Check: Do you already have the assets needed for recovery?

Look for a cluster of related pages ranking in positions 6–20. If it exists, the conditions for Archetype 1 may apply. If it doesn't, the structural case for content-driven recovery is weak. Continue to Chapter 8.

And if headline traffic appears to be recovering mainly through branded queries, you're seeing product momentum rather than content recovery.



WORTH AUDITING

Before producing new content, audit existing pages with strong backlink profiles ranking in positions 8–25. Those pages may be able to break through without new production.

Before you move on

Use these three questions to determine whether a recovery archetype applies before committing to a plan.

QUESTION	WHAT TO CHECK	NEXT STEP
Where did losses concentrate?	Did losses hit peripheral content or foundational pages?	Foundational → No archetype applies. Redirect investment. Peripheral → Go to step 2.
Did positions 1-3 hold?	Did top-position pages hold throughout the decline?	Held → Assess Archetype 2. Fell → Continue to Step 3
Do you have the assets needed for recovery?	A cluster of related pages or an underperforming high-equity page?	Cluster → Assess Archetype 1 Neither → Continue to Chapter 8



THE BOTTOM LINE

Before asking how to recover, ask whether recovery is still possible. If the structural floor survived, the opportunity may already exist in your content. If it didn't, the recovery playbook changes.

CHAPTER

08

When Recovery Isn’t the Right Strategy

Some declines are not recovery problems. They are structural failures.

The recovery archetypes in the previous chapter only applied when the structural foundation remained intact. The cohort also revealed situations where that foundation had already broken.

In these cases, the issue wasn’t lack of effort. More content, more links, or more refreshes were never going to change the outcome.

This chapter looks at the cases where losses had reached the foundation of the program. Once that structural floor is gone, standard organic investment no longer matches the problem.

IN THIS REPORT

01 — What Happened to B2B Traffic

02 — Assumptions That Didn’t Survive the Data

03 — What Your Traffic Number Hides

04 — Which Force Hit You

05 — The Foundation of Durable Demand

06 — When Growth Starts to Crack

07 — How Companies Actually Recovered

08 — When Recovery Isn’t the Right Strategy

09 — Understanding AI Visibility

10 — Building for AI Visibility

11 — The Durable Demand Playbook

What structural decline looks like

Three patterns appeared consistently across the non-recovering companies in this study. In each, standard SEO responses like publishing more content, earning more links, or refreshing existing pages were unlikely to produce recovery.

Tactical declines can respond to SEO investment because the underlying problem is still fixable through content or authority. Structural declines usually can’t because the foundation that once produced growth no longer exists. The distinction determines whether the right move is to invest or redirect.

THE NEXT QUESTION

If recovery isn’t guaranteed, what makes one decline recoverable and another permanent? The answer begins with one distinction: tactical versus structural decline.

Pillar content collapse



WHAT HAPPENED

Pages that defined the program's core topical authority lost their rankings. These were the foundational pages that established the site's authority in its primary topic.



WHY RECOVERY DIDN'T FOLLOW

Pillar pages anchor the rest of the program. When they fall, the supporting content loses the structural foundation it depended on. Publishing more peripheral content won't rebuild it.



WHAT TO CHECK

Pull ranking history for the three to five pages most central to your topical authority. These are the pages that define your primary category claim in organic search.

If those pages lost significant rankings and haven't recovered, the preconditions for the Chapter 7 recovery archetypes may no longer exist.



KEY INSIGHT

When pillar pages collapse, you're looking at a rebuild project instead of a repair project.

Architectural damage



WHAT HAPPENED

A redesign or page overhaul broke existing rankings even though the URLs and backlink equity remained intact. The pages changed enough that the rankings they had earned no longer held.



WHY RECOVERY DIDN'T FOLLOW

The problem wasn't authority or content volume. The redesign changed the on-page signals that had earned the rankings, so publishing new content elsewhere couldn't fix it. The fix was architectural, not editorial.



WHAT TO CHECK

If a significant traffic decline aligns closely with a redesign, platform migration, CMS change, or major content overhaul, architectural damage is the leading explanation.



KEY INSIGHT

Publishing more content won't repair rankings lost by a broken page. Restore the page before expanding the program.

Upstream demand substitution



WHAT HAPPENED

Search volume for the affected queries declined. The queries themselves occurred less often, most likely because users got the answer directly in AI tools instead of searching for it.



WHY RECOVERY DIDN'T FOLLOW

This loss happens before the search. If users generate the answer directly in an AI tool, the query never occurs. There is no ranking to defend and no content refresh that can bring that demand back.



WHAT TO CHECK

Pull search volume trends for the queries tied to your declining pages. If search volume has fallen over 12–18 months on the same queries where traffic declined, and rankings or CTR don't explain the drop, upstream demand substitution is the leading explanation. The question is no longer "How do we recover traffic?" It's "Where does durable demand still exist?"



KEY INSIGHT

You can't recover demand that no longer reaches a search engine. The opportunity is to find where that demand has moved.

The structural vs. tactical test

QUESTION 1: Are the losses on foundational or peripheral content?

Map the declined pages against their role in the program. If the losses include your highest-authority pages, the ones that define your topical claim, the decline may be structural. If losses are concentrated in peripheral content while core pages held, the decline is more likely tactical and the Chapter 7 archetypes may still apply.

QUESTION 2: Did the decline align with a specific internal event?

If the drop aligns closely with a redesign, migration, CMS change, or major content overhaul, architectural damage is the leading explanation. Check whether referring domains held while rankings fell. That's the signature.

QUESTION 3: Is search volume declining?

Pull 12–18 months of search volume data for the queries tied to the decline. If search volume is falling independent of rankings, demand is likely shifting upstream. No ranking improvement or content investment can restore demand that no longer shows up as search behavior.

If one signal is strong, pause before investing. If two or more are present, treat the decline as structural until proven otherwise. At that point, recovery efforts are more likely to create cost than recovery.



DECISION RULE

Don't choose a recovery strategy until you've diagnosed the type of decline. The wrong diagnosis often leads to months of work that can't solve the underlying problem.

FIGURE 8.1

Structural vs. Tactical Decision Matrix

Match your observed signal to a row before spending anything. Structural patterns call for redirection, not reinvestment.

PATTERN	PRIMARY EVIDENCE SIGNAL	WHAT STANDARD SEO INVESTMENT PRODUCES	RESPONSE
STRUCTURAL PATTERNS			
Pillar content collapse	Majority of tracked keywords lost on one or two cornerstone pages. Cluster pages hold steady.	Content refresh cycles on the affected page show no ranking or traffic response.	Redirect
Architectural damage	Referring domains held. Rankings dropped despite an intact link profile.	Continued content and link-building spend shows no traffic recovery because the root cause is untouched.	Redirect
Demand substitution	Cluster-level impressions decline in GSC even as positions remain unchanged.	Optimization work has no effect. There is no addressable demand left to capture.	Redirect
TACTICAL PATTERNS			
Algorithm update damage	Traffic drop aligns to a known update date, and positions partially recovered in subsequent months.	Standard content-quality and technical fixes typically restore position within one to two update cycles.	Recover
Competitive displacement	Rankings drop while a specific competitor's pages rise into the vacated positions. Referring domains and content quality are unchanged.	Closing the specific content or link-quality gap typically regains the lost positions.	Recover

Examples from the cohort

Secureframe & Sprinto (pillar content collapse)

Secureframe and Sprinto both lost foundational compliance pages like “What Is SOC 2?” and “What Is ISO 27001?” These weren’t supporting articles. They were the pages that established each site’s authority in compliance, and neither recovered during the study window.

Hyperproof provides the contrast. It lost peripheral compliance content while its core GRC pages held, preserving the structural foundation that later enabled recovery. Secureframe and Sprinto lost that foundation, removing the precondition for a Chapter 7 recovery.

Freshworks (architectural damage)

Freshworks’ December 2025 Dew design system launch aligned closely with a broad ranking decline. Referring domains remained stable while rankings fell, pointing to architectural damage rather than a loss of authority.

The recovery problem wasn’t a lack of content. It was that the pages no longer sent the signals that had earned their rankings.

Aha! (upstream demand substitution)

Aha! lost 60.3% of traffic across 154 template pages. The clearest signal wasn’t rankings or CTR. It was declining search volume for the affected queries, suggesting users increasingly generated templates directly in AI tools instead of searching for them.

The problem wasn’t that Aha! stopped ranking. It was that fewer searches were happening in the first place.



KEY INSIGHT

Once the underlying failure mechanism was identified, the recovery outcome became far more predictable.

Before you move on

Run these three checks before committing recovery budget.

SIGNAL	HOW TO CHECK	DIAGNOSIS
Pillar content losses	Did your 3–5 foundational pages lose rankings and fail to recover?	Chapter 7 recovery archetypes are unlikely to apply. Redirect before investing.
Redesign-aligned timing	Did the decline begin after a redesign, migration, or CMS change? Did referring domains hold while rankings fell?	Architectural damage likely. Fix the architecture before publishing anything new.
Declining search volume	Is search volume falling independent of rankings?	Upstream demand substitution. Recovery won't restore demand.



THE BOTTOM LINE

More SEO investment won't produce recovery when the decline is structural. In most cases, you can identify that before spending anything.

CHAPTER

09

Understanding AI Visibility

By mid-2025, the dominant view in B2B marketing was that AI Overviews threatened informational content. If Google could answer the query in the SERP, why would anyone click through to the page?

The cohort supports the concern, but not the conclusion. AI Overview suppression is real and sometimes severe, but it doesn't affect informational content uniformly.

AI Overviews affect informational queries that no longer require a click.

This chapter examines three distinct outcomes: resistant content, direct SERP interception, and complete query routing to specialist tools. The right response depends on which one you're seeing.

IN THIS REPORT

01 — What Happened to B2B Traffic

02 — Assumptions That Didn't Survive the Data

03 — What Your Traffic Number Hides

04 — Which Force Hit You

05 — The Foundation of Durable Demand

06 — When Growth Starts to Crack

07 — How Companies Actually Recovered

08 — When Recovery Isn't the Right Strategy

09 — Understanding AI Visibility

10 — Building for AI Visibility

11 — The Durable Demand Playbook

The resistance evidence: 5 categories, 1 pattern

Before looking at suppression, we should be precise about what resisted it.

Five categories showed definitional content holding or growing throughout the AI Overview expansion window, the same period when industry commentary predicted broad suppression of educational content.

These weren't niche programs in unusual conditions. They competed in crowded categories where well-funded peers were publishing similar educational content. Yet their definitional pages held up and, in several cases, continued to grow while AI Overviews expanded.

That doesn't prove AI Overviews have no effect on definitional content. It supports a narrower and more useful conclusion: during this period, the effect wasn't broad or severe enough to appear consistently in traffic data. Seeing the same pattern across five companies in five different categories is difficult to dismiss as noise.



KEY INSIGHT

If AI Overviews broadly suppressed educational content, our categories should have declined together. They didn't.

The direct evidence: When rankings hold but traffic disappears

The resistance evidence showed what AI Overviews didn't suppress. Aha! provides the clearest example of what they did.

Across more than 20 queries, rankings held or improved while traffic fell between 79% and 99%. That isn't the pattern you'd expect from competitive displacement, where rankings and traffic typically move together.

The next figure shows five representative examples.

WHAT THE DATA SHOWED

Across more than 20 queries, rankings held or improved while traffic fell 79–99%. In this context, the pattern is most consistent with AI Overview interception rather than competitive displacement.

Competitive displacement changes rankings first. AI Overview interception changes clicks first.

FIGURE 9.1

Rankings Held. Traffic Cratered.

Ahrefs keyword data, sorted by severity. Rankings held or improved on every query while traffic fell 67–93%, pointing to SERP interception rather than competitive displacement.

KEYWORD	RANKING MOVEMENT	PREV. TRAFFIC	CURR. TRAFFIC	CHANGE
creative brief template	2 → 1	700	46	-93%
product roadmap template	3 → 1	357	53	-85%
product roadmap	5 → 2	623	167	-73%
product development	1 → 1	867	273	-69%
competitive analysis template	3 → 1	374	122	-67%

Source: Ahrefs keyword export, Aha! domain. Comparison period July 2024 vs. July 2026.

What determines exposure

Aha!'s keyword data doesn't show random loss. The queries with the sharpest CTR collapse all share the same characteristic: the user can complete the task in the SERP without clicking.

- "Marketing job titles" can be answered with a list.
- "User story template" can be supplied directly.
- "Product marketing" can be defined in a short Overview.
- "Product positioning" can be summarized into a simple framework.

In each case, the SERP satisfies the user's need before they ever reach the ranking page.

This suggests AI Overview exposure is driven less by whether content is informational than by whether the query can be resolved without a click.

The next figure illustrates that spectrum.

KEY INSIGHT

The evidence points to a different framework than the industry narrative. AI Overviews don't suppress informational content as a category. They suppress queries that can be fully satisfied in the SERP. Where a query falls on that spectrum is a stronger predictor of exposure than its topic or format.

FIGURE 9.2

The click dependency spectrum

Where a query sits on this spectrum, not its topic or format, determines whether it can be intercepted in the SERP.



NO CLICK NEEDED

- Role & career lists
- Templates
- Shallow term definitions
- Step-by-step routine processes

MIXED

- Strategic overviews
- Comparison guides
- Multi-factor explanations

ENGAGEMENT-REQUIRING

- Comprehensive methodology guides
- Proprietary research
- Deep conceptual frameworks

The extreme case: Complete query routing

Aha! shows AI Overview interception: the page still ranks, but the click is absorbed in the SERP.

The study also surfaced a more extreme pattern.

During the AI citation phase, Freshworks's "is this website down" query was tested across six AI engines. Every engine routed users to dedicated uptime tools such as Downdetector, IsItDownRightNow, and UptimeRobot. None cited Freshworks. None answered from editorial content.

This is fundamentally different from AI Overview interception. With interception, the page still has an opportunity to earn the click. With complete query routing, content never enters consideration because the AI determines the query requires a live data source instead.

For operational queries such as uptime checks, live pricing, availability, and current status, that appears to be a structural shift rather than a temporary one. No editorial page, regardless of quality, can substitute for live data.



KEY INSIGHT

Complete query routing is distinct from AI Overview interception. Once an AI system determines a query requires a specialist tool rather than editorial content, improving the page cannot restore visibility because the content is no longer competing in the same layer.

Before you move on

1 Inspect the SERP.

Search the primary keyword in an incognito window and document the full results page. Look beyond rankings. Note AI Overviews, featured snippets, citations, shopping results, videos, and any other feature that could intercept the click.

2 Ask one question.

Can the user complete the task from the SERP alone, or do they still need to click?

Lists, templates, brief definitions, and routine processes are the clearest exposure patterns. Comprehensive guides, proprietary frameworks, and methodology content are generally more resistant. If you're unsure, use Figure 9.2 as your reference.

3 Prioritize by opportunity.

Start with pages combining high traffic and high satisfiability risk.

The response is usually one of three actions:

- Deepen the page with proprietary research or methodology.
- Reframe the topic toward a question requiring engagement.
- Deprioritize content structurally easy to satisfy in the SERP.



THE BOTTOM LINE

AI Overviews don't suppress informational content uniformly. They suppress queries the SERP can satisfy without a click. The risk is specific, auditable, and unevenly distributed across most content programs.

CHAPTER

10

Building for AI Visibility

Two competing narratives have shaped the debate around AI and content discovery. One argues that AI systems require an entirely separate content strategy. The other assumes strong SEO automatically produces AI visibility.

The evidence from this study supports neither.

Content that earns durable organic authority often earns AI citations, but not always. The difference is usually driven by query type, not content quality. That distinction determines whether an AI visibility gap is fixable or whether the AI is routing the query where editorial content cannot compete.

The question is no longer whether your page ranks. It's whether the AI system finds and cites it at all.

IN THIS REPORT

01 — What Happened to B2B Traffic

02 — Assumptions That Didn't Survive the Data

03 — What Your Traffic Number Hides

04 — Which Force Hit You

05 — The Foundation of Durable Demand

06 — When Growth Starts to Crack

07 — How Companies Actually Recovered

08 — When Recovery Isn't the Right Strategy

09 — Understanding AI Visibility

10 — Building for AI Visibility

11 — The Durable Demand Playbook

What the citation data shows

This study tested 13 prompts across six AI engines: ChatGPT, Copilot, Google AI Mode, Google AI Overviews, Gemini, and Perplexity. Each prompt targeted a query central to one company's organic growth story, covering 13 of the 15 deep-dive companies.

Across 78 engine-prompt combinations, AI systems cited the target company's domain in 9 of the 13 cases. In 7 of those cases, the engines cited the exact target URL. In the remaining 2, they cited a different page on the same domain. Four companies received no citations from any engine.

This validation set is intentionally small and does not represent a universal model of AI retrieval. Even so, it reveals a clear pattern: AI systems usually cited companies that had already established strong organic authority.

Blog and educational content accounted for 5 of the 9 confirmed citation cases. Glossary and definitional pages accounted for 2. Product and marketing pages accounted for the fewest. These results mirror the durability findings from Chapter 5. The content that proved most resilient in organic search was also the content AI systems most often retrieved.



KEY INSIGHT

AI systems cited the target company in 9 of 13 tested cases, usually pointing to the exact page that had built that company's organic authority. Blog and educational content earned the most citations, reinforcing the broader pattern that the content most resilient in search is often the same content AI systems retrieve.

What AI systems tend to cite

The content AI systems cited most consistently across engines shared three observable properties.

SPECIFICITY

The strongest-performing pages answered one specific, well-defined question instead of a broad overview. CircleCI's What Is Continuous Integration? page and Rippling's HR glossary entries each focused on a single concept and answered it directly. Broad introductory content was cited less often.

ORIGINAL POV

The most-cited pages offered something beyond widely available information. They offered a proprietary definition, a domain-specific explanation, or a distinctive framework instead of simply restating common knowledge.

DIRECT QUESTION MATCHING

The strongest pages answered the exact question users asked, clearly and completely. That sounds obvious, but it has an important implication: pages built to answer one question thoroughly appeared more retrievable than pages optimized primarily for breadth, keyword coverage, or topical authority.



KEY INSIGHT

These are also the qualities that support durable organic search performance. The retrieval mechanisms differ, but the overlap is substantial. The page that earns search authority by directly answering a question is often the same page an AI system retrieves when asked that question.

Two structural exceptions

Several non-cited cases in the cohort were not content problems at all. Instead, the query itself routed AI systems away from editorial content. Two patterns appeared most often.

OPERATIONAL QUERIES: Routing to specialist tools

Freshworks' tested query asked whether its service was currently operational. Across all six AI engines, the responses cited multiple uptime monitoring tools. No engine cited Freshworks or any other editorial content.

This matches the complete query routing pattern identified in Chapter 9. AI systems treated the query as a request for live operational data, not an explanation. The limitation wasn't Freshworks' content. It was the query itself. No editorial strategy changes that routing decision.

INSTITUTIONAL AUTHORITY QUERIE: Routing to primary sources

Vanta's tested query focused on risk management frameworks, the same category as its largest traffic-gaining page during the study window. Yet AI systems consistently cited primary authorities such as NIST, AWS, and other regulatory sources instead of Vanta's explanation.

This is a different form of structural routing. AI systems could answer the query with editorial content, but they consistently preferred the original source material over brand content built around it. The limitation wasn't content quality. It was the query's preference for institutional authority.

Not every missing AI citation reflects a content quality problem. When AI systems route queries to specialist tools or primary authorities, the outcome is structural. Query type—not content strength—determines who gets cited.



KEY INSIGHT

These patterns also change how AI visibility should be measured. If a query consistently routes to specialist tools or primary sources, it isn't a useful AI discoverability benchmark. Exclude it from your KPIs and focus on queries where editorial content competes.

Engine behavior differs

Not all AI engines retrieved the same content or cited the same companies at the same rate. In this validation set, **Google AI Mode** and Google **AI Overviews** cited cohort content most frequently, while **Perplexity** drew from a broader mix of sources. **ChatGPT** and **Copilot** cited the fewest cohort pages, often responding without citing a specific source.

These differences matter for measurement. AI discoverability is not a single channel but a collection of retrieval systems with different citation behaviors. Measuring visibility through one engine alone can produce a very different picture than measuring across multiple engines.

FIGURE 10.1

Citation Rates Differed Sharply by Engine

For each company, six AI engines were asked the prompt equivalent of its core growth query. Cells show whether the engine cited the exact target page, another URL on the same domain, or nothing at all.

Company	ChatGPT	Copilot	Google AI Mode	Google AI Overviews	Google Gemini	Perplexity	Total
CircleCI	○	○	●	●	●	●	4/6
Braze	○	●	●	○	○	○	2/6
CloudBees	○	○	○	●	○	●	2/6
Harness	○	●	○	●	○	○	2/6
Navan	○	○	●	●	○	○	2/6
PostHog	○	○	●	●	○	○	2/6
Rippling	○	○	●	●	○	○	2/6
Ramp	○	○	○	●	○	○	1/6
ZoomInfo	○	○	◐	○	○	○	1/6
Akita	○	○	○	○	○	○	0/6
Freshworks	○	○	○	○	○	○	0/6
Gainsight	○	○	○	○	○	○	0/6
Vanta	○	○	○	○	○	○	0/6
Engine Total	**0**	**2**	**6**	**7**	**1**	**2**	

● Cited at exact target URL ◐ Cited at different URL on same domain ○ Not cited

Before you move on

Test your three most important category queries across the AI engines your audience actually uses. Focus on queries where editorial content should still compete.

For each query, ask:

- Is your domain cited?
- If not, who is?
- Does the query consistently route to specialist tools or primary sources?

If a competitor is cited and you are not, the gap is probably fixable through better content. If every engine routes the query away from editorial content, stop measuring AI visibility against that query.



THE BOTTOM LINE

Content with durable organic authority often earns AI citations. But when AI systems route a query to specialist tools or primary sources, the gap is structural rather than editorial. In those cases, improving the content won't change the outcome.
Before you move on

CHAPTER

11

The Durable Demand Playbook

Every chapter in this report has been a diagnosis. This one is the prescription.

Across the study, the same investment question kept resurfacing: does content build durable demand, or does it borrow fragile traffic from a system that no longer rewards it? This chapter turns the study's findings into a practical framework for deciding what to build next.

Content type was the strongest pattern observed. Specifically, the value the content created, whether it required engagement, and whether the user's goal required a click. Strong companies invested in content with those characteristics. Companies that declined depended more heavily on content without them.

What follows is a five-step framework distilled from the patterns that most consistently separated durable growth from fragile traffic across the cohort.

IN THIS REPORT

01 — What Happened to B2B Traffic

02 — Assumptions That Didn't Survive the Data

03 — What Your Traffic Number Hides

04 — Which Force Hit You

05 — The Foundation of Durable Demand

06 — When Growth Starts to Crack

07 — How Companies Actually Recovered

08 — When Recovery Isn't the Right Strategy

09 — Understanding AI Visibility

10 — Building for AI Visibility

11 — The Durable Demand Playbook

Putting the evidence to work

Every major finding in this report points back to the same conclusion: investing in the right type of content drives sustainable organic growth far more than simply publishing more content.

The Durable Demand Framework distills those findings into five principles. The sections that follow explain each principle and the evidence behind it.

FIGURE 11.1



A closer look at the framework

1. DIAGNOSE THE MECHANISM

Before investing in recovery or growth, identify what actually drives the traffic pattern. Traffic data shows what happened; position, timing, and search volume explain why. The wrong diagnosis leads to the wrong investment.

2. BUILD ENGAGEMENT-REQUIRING ASSETS

Invest in content whose value requires engagement, not content that can be satisfied in a SERP or AI response. Across this study, that was the defining difference between durable and fragile demand.

3. BUILD A PORTFOLIO OF DURABLE CONTENT TYPES

Avoid overreliance on any single page or content type. Programs that compound distribute growth across multiple durable content types rather than concentrating it in one asset.

4. MEASURE PORTFOLIO HEALTH

Track cannibalization, position consolidation, branded/non-branded split, and single-page concentration. These signals reveal whether a program compounds authority or simply replaces losses.

5. EVALUATE EVERY INVESTMENT WITH TWO QUESTIONS

Before committing to any content investment, ask two questions: Does this require engagement to deliver its value, or can it be extracted in a single pass? Can AI fully answer the query without a click?



KEY INSIGHT

Durable demand comes from content that earns engagement, not content that simply answers a question. Every recommendation in this playbook follows from that principle.

FIGURE 11.2

The Durable Demand Matrix

Four content tiers, scored against AI Overview vulnerability, AI citation likelihood, and observed durability across the study cohort.

CONTENT CATEGORY	AIO VULNERABILITY	AI CITATION LIKELIHOOD	HISTORICAL DURABILITY	ACTION IMPLICATION
Foundational learning-demand content	Medium	High	Durable	Invest at depth.
In-depth educational and strategic content	Medium	High	Conditional	Invest selectively.
Proprietary methodology / original research	Low	High	Durable	Core investment.
Template, list, basic how-to, actionable output	High	Low	Fragile	De-invest or transform.

Synthesizes three evidence streams: cohort traffic behavior across 118 tracked companies, AI Overview / CTR analysis (Ch. 9), and AI citation testing across 78 engine-prompt combinations.

The Durable Demand Framework explains how to evaluate content investments. **This matrix translates those principles into practical content classification tiers observed across the study.** Rather than grouping pages by format alone, it organizes them by observed durability, AI vulnerability, and structural defensibility. The next section examines each tier in detail and explains why it behaved the way it did.

Durable: Foundational learning demand

Foundational learning demand comes from questions every new practitioner, team, or organization must answer before they can move forward. Unlike trend-driven topics, these concepts remain relevant because the audience continually renews. New professionals enter the field, organizations adopt unfamiliar technologies, and teams encounter the same foundational questions for the first time.

This continual renewal makes foundational learning one of the most durable forms of content. The concept itself may not change, but demand persists because new audiences keep arriving. Organizations that invest in strong foundational resources are building assets that continue attracting new learners long after publication.



KEY INSIGHT

Foundational demand is durable because the audience continually renews. The concept does not need to be new for demand to remain durable.

Durable with conditions: Engagement-requiring educational content

Not all educational content is equally durable. Its long-term value depends less on the topic than on how the user consumes it. When the content requires readers to think, compare, analyze, or apply information, it remains difficult to replace with a single AI response or search result.

That condition matters more than the format itself. A methodology guide, comparison page, or strategic playbook can become a durable asset when its value comes from sustained engagement. If the useful information can be extracted into a short answer or checklist, however, the same content becomes structurally fragile.



KEY INSIGHT

Educational content is durable only when completing the user's goal requires engagement. If AI can satisfy the query without a click, the investment becomes structurally fragile.

CHAPTER 11 — The Third Durability Tier

Structurally protected: Proprietary methodology & original research

The most defensible content is built on knowledge that competitors and AI systems cannot reproduce because it originates with you. Original research, proprietary frameworks, unique methodologies, and firsthand data create value by contributing something new rather than summarizing what already exists.

These assets are naturally resistant to AI substitution. While AI can explain or reference original work, it cannot independently generate the underlying evidence or experience that gives it authority. That makes proprietary knowledge one of the strongest long-term investments for organizations seeking durable organic demand.



KEY INSIGHT

The strongest competitive advantage isn't explaining existing knowledge. It's creating knowledge that only your organization can provide.

Fragile: Extractable output content

Extractable output content delivers its primary value immediately as an answer, list, template, or sequence of steps. When users can accomplish their goal without engaging with the page itself, the traffic supporting that content becomes structurally vulnerable.

This does not mean these assets lack value. Many continue to rank well and serve important business functions. However, because their core utility can often be reproduced directly within search results or AI-generated responses, they face greater long-term pressure than content that requires sustained engagement or contributes original knowledge.



KEY INSIGHT

When the user's goal can be completed without visiting the page, the traffic supporting that content becomes structurally fragile.

The two-question test

Every content investment can be evaluated with two questions:

1. DOES THIS CREATE ONGOING ENGAGEMENT?

Will readers return to this content as their situation changes? Will they share it? Does understanding it require reading, thinking, or applying, or can the useful part be extracted in a single pass?

Content that requires ongoing engagement will often compound. Content that can be consumed once and forgotten rarely does.

2. CAN AI FULLY ANSWER THE USER'S QUERY?

Search the primary query in an incognito window. If an AI Overview appears, ask whether the user can accomplish their goal from the answer alone. If they still need to click, the content is durable. If they don't, the traffic supporting that content is inherently more vulnerable, regardless of ranking position or execution quality.



KEY INSIGHT

If the answer to Question 1 is yes and the answer to Question 2 is no, you have the foundation of a durable content asset.

FIGURE 11.3

Portfolio Health Scorecard

Five diagnostics across three families, benchmarked against the winning programs in the study cohort.

METRIC & HOW TO CALCULATE	HEALTHY SIGNAL	INVESTIGATE WHEN
GROWTH QUALITY		
Cannibalization rate Gross traffic loss / gross traffic gain	< 30%	> 60%
Single-page concentration Percent of total gross gains driven by the largest gaining page	Top page < 20%	Top page dominates gains
TRAFFIC COMPOSITION		
Branded / non-branded split Percent of total gross gains driven by the largest gaining page	Non-branded primary	Branded > 40% of growth
PORTFOLIO DURABILITY		
Position distribution Share of tracked keywords ranking in positions 1-10	1–10 share rising	1–10 share declining
Content type mix Share of traffic held by Tier 1–3 content	Tiers 1–3 majority	Tier 4 majority

The Durable Demand Test helps evaluate individual pages. This scorecard applies the same principles across an entire content portfolio. Across the highest-performing programs in this study, these five diagnostics appeared consistently. Together, they provide a practical framework for evaluating whether a content program compounds authority or simply replaces yesterday's gains.

Investment decisions

Every content investment should fall into one of three buckets. In doing so, you'll invest more deliberately in content with durable demand.



DO NOT INVEST

Default action: Stop investing in SEO.

- Template libraries whose value is the downloadable asset
- Role, career, and title lists
- Basic how-to content for mature, well-documented processes

Reason: The user's goal can be completed without visiting the page.



DE-RISK

Default action: Keep, but transform.

- Educational content lacking proprietary depth
- Clusters with satisfiable queries
- Assets succeeding for reasons you don't yet understand

Reason: Increase engagement, original insight, or methodology before demand erodes.



DOUBLE DOWN

Default action: Increase investment.

- Foundational learning content
- Proprietary research & methodology
- Interactive tools and calculators

Reason: These assets showed the strongest structural durability across the study.



KEY INSIGHT

Don't allocate resources based on content format. Allocate them based on structural durability.

Apply the playbook

By the end of these four steps, you'll have a clear picture of which parts of your content program are durable, which are vulnerable, and where future investment belongs.

1. CLASSIFY YOUR TOP 50 PAGES

Use the Durable Demand Matrix (Figure 11.2) to classify your highest-traffic pages by durability tier. This reveals how much of your traffic depends on structurally fragile content.

2. MEASURE PORTFOLIO HEALTH

Complete the five diagnostics in the Portfolio Health Scorecard (Figure 11.3) to establish your baseline. Focus on cannibalization, single-page concentration, branded/non-branded growth, position distribution, and content mix.

3. PRESSURE-TEST YOUR HIGHEST-TRAFFIC PAGES

Search the primary query for your top pages in an incognito window and apply the Two-Question Test. If users can complete their goal without clicking, prioritize those pages for review.

4. AUDIT YOUR ROADMAP

Run every planned content investment through the Two-Question Test before work begins. If it cannot be restructured to require engagement, reconsider whether it deserves investment.

5. EVALUATE EVERY INVESTMENT WITH TWO QUESTIONS

Make the Two-Question Test a standing check on all future content decisions, not just the initial roadmap audit. Before work begins on any piece: Does it require engagement to deliver value? Can AI fully satisfy the query without a click?



THE BOTTOM LINE

You should now know three things: what to stop producing, what to restructure, and where to invest more aggressively. More importantly, you'll have a framework for evaluating future content decisions for durable demand rather than historical SEO assumptions.

WRAP-UP

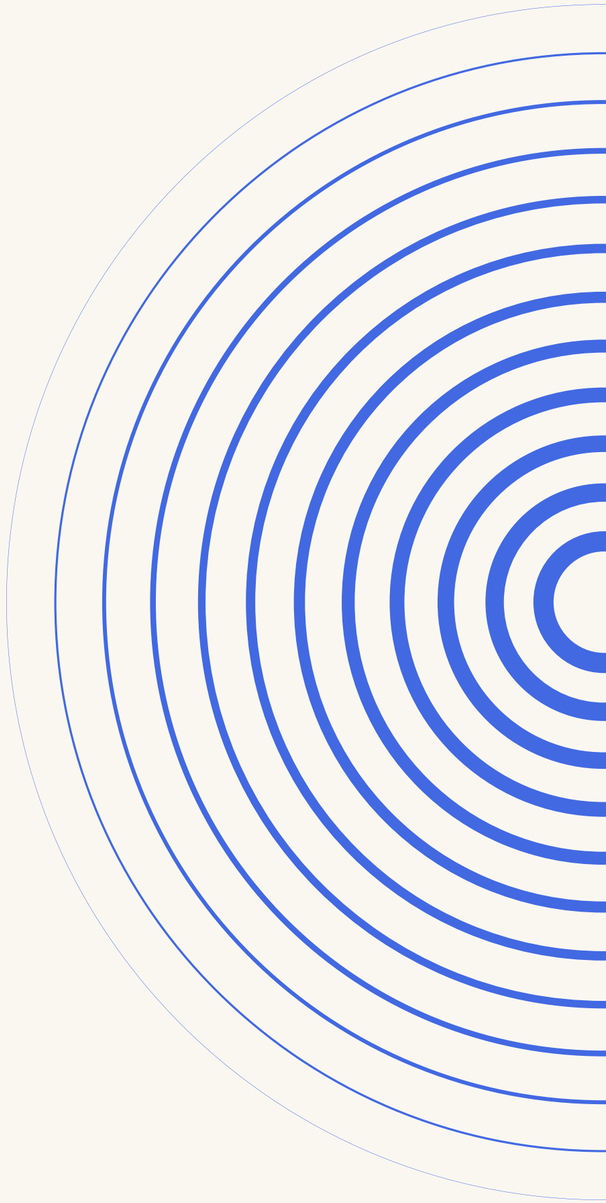
Search didn’t die. Fragile demand did.

The companies that continued growing in 2025 and 2026 weren’t lucky, and they weren’t early. They built content that required something from the reader rather than content that provided a quick answer and lost the user the moment the answer was given.

That is the difference between chasing traffic and building demand. Traffic can be intercepted. Demand can’t, at least not when it renews because the problem keeps recurring and the audience keeps arriving.

The organic landscape will keep shifting. AI capabilities will advance. Search behavior will continue to evolve. The programs that survive those shifts will be the ones that answer two questions correctly: Does this require engagement? Can AI satisfy it without a click?

The future of organic growth won’t be won by producing more answers. It will be won by creating reasons to return. In an internet increasingly optimized to answer questions, enduring advantage belongs to the brands people still need after they have the answer.



Appendix A

Research Methodology

This appendix documents how the 118-company cohort was selected, measured, classified, and analyzed. It also summarizes the study's quality controls. Metric formulas appear in Appendix C.

Research Objective

This study examined one question:
What distinguished B2B SaaS companies that continued growing organic traffic during a period of expanding AI visibility and algorithm volatility from those that stalled or declined?
The research used an observational longitudinal design across an 18-month study window. It identifies recurring associations rather than experimental causality. The consistency of the findings across companies, categories, and analytical methods provides strong support for the report's directional conclusions.

Cohort Design and Selection

The study used a category-stratified cohort of B2B SaaS companies with sufficient operating and organic search history to support an 18-month analysis.

Approximately 20–30 candidates were identified in each category before screening. Candidates were sourced primarily through G2 category pages, with Crozdesk, GetApp, and Capterra used when additional candidates were needed.

Ten categories

- Sales Engagement / Sales Intelligence
- Marketing Automation
- Customer Support / Help Desk
- Customer Success
- HR / People Operations
- Finance / Spend Management
- Project Management / Collaboration
- Developer Tools / DevOps
- Security / Compliance
- Product Analytics / Product Management

Inclusion gates

A company was included only when it met all five requirements:

1. Its primary business was a B2B SaaS product.
2. It had an independently interpretable operating history.
3. Its primary domain and operations remained sufficiently continuous during the study window.
4. Its core product clearly fit the category being studied.
5. Ahrefs reported trackable organic traffic at both ends of the study window.

Employee count was recorded but not used for inclusion.

Every inclusion and exclusion decision was preserved in the screening archive with the gate and rationale that determined it.

Companies affected by acquisitions, domain migrations, operating discontinuity, or ambiguous category fit were excluded when those conditions prevented a defensible comparison.

Three analytical phases

- Phase 1: Candidate screening and cohort selection
- Phase 2: Full-cohort measurement and traffic classification
- Phase 3: Company deep dives, comparison briefs, and AI citation validation

Measurement and Traffic Classification

Organic performance was measured across the 18-month window from January 2025 through June 2026. Three-month averages were used at both ends of the window to reduce monthly volatility. All reported changes use these averages unless

Appendix A

Research Methodology

otherwise noted.

Traffic change = (end-period average – start-period average) : start-period average × 100
Start period: January–March 2025 End period: April–June 2026

Primary traffic classifications

Grew — > +10%
Flat — –10% to +10%
Declined — < –10%

Thresholds were locked before classification began and applied uniformly across all 118 companies. Manual review was used only when a near-zero starting value made percentage change misleading. Every override was documented.

Final distribution: 44 grew · 19 remained flat · 55 declined

Full cohort fields

For every company, the study recorded:

Performance

- Organic traffic and change
- Traffic value and change
- Referring domains and change

Peak context

- Study-window peak and month
- Historical peak and month
- Distance from historical peak

Classification

- Grew / Flat / Declined
- Traffic trajectory

Phase 2 used only standardized company-level Ahrefs measurements. Page-, query-, and position-level exports were introduced during Phase 3.

T**raffic trajectory taxonomy**

Trajectory labels describe traffic shape only. They do not affect a company's Grew, Flat, or Declined classification.

Nine trajectory labels were used: Continuous Acceleration, Steady Growth, Growth then Pullback, Early Peak then Plateau, Early Peak then Decline, Decline then Recovery, Flat, Steady Decline, and Volatile.

Deep-Dive Analysis and Mechanism Attribution

Fifteen companies received structured Phase 3 deep dives. Selection was designed to cover five recurring questions:

- How do successful programs continue accelerating?
- Why do previously successful programs stall or decline?
- How do companies recover after sustained decline?
- Why do some companies fail to recover?
- How do companies break out of prolonged stagnation?

The set deliberately included growing, flat, and declining companies. Declining companies were included to distinguish durable growth from conditions that only appeared healthy.

Standardized analysis

Every deep dive used the same 14-field framework covering performance data, historical context, page-level analysis, authority signals, algorithm and AI timing, competing explanations, and final mechanism attribution.
Each analysis included an explicit competing-explanations step before a mechanism was assigned. Funding events, product launches, category shifts,

Appendix A

Research Methodology

acquisitions, domain changes, competitor exits, and other off-content explanations were evaluated rather than assumed away.

Standardized external timelines

Verified timelines of major Google algorithm updates and AI search expansion events were compiled before company-level synthesis and applied consistently across all deep dives. This allowed the study to compare company inflection points against external events without retroactively selecting dates to fit individual cases.

Comparison briefs

Nine comparison briefs examined companies facing similar category conditions but producing divergent outcomes. Each followed the same structure: shared context → divergent trajectory → mechanism comparison → cross-company conclusion. These comparisons strengthened the analysis by showing which differences persisted among companies operating in similar category conditions.

Content Growth-Pattern Classification and AI Citation Testing

Content growth-pattern classification

After the deep dives, all 44 growing companies were reviewed to identify the content pattern contributing most strongly to their observed organic growth during the study window.

Five primary patterns were used: definitional and educational content, brand and product demand, two confirmed co-primary patterns, programmatic content, and interactive tools and free resources.

Deep-dive companies were classified from the full Phase 3 evidence. Remaining companies were assessed through site architecture, representative pages, and Ahrefs Top Pages exports where needed. Classification was assigned to the content pattern that contributed most meaningfully to both traffic composition and growth during the study window.

Of the 44 growing companies: 38 were classifiable at full weight; 2 were classifiable but treated as low-weight evidence due to limited traffic scale; and 4 were excluded from pattern synthesis because technical artifacts or unit-of-analysis problems prevented a defensible classification. The distribution reported in the main study uses the 38 full-weight cases

AI citation validation

AI citation behavior was tested for 13 of the 15 deep-dive companies using Scrunch across six engines: ChatGPT, Copilot, Google AI Mode, Google AI Overviews, Gemini, and Perplexity. Hyperproof and LogRocket were added to the deep-dive cohort after citation testing was complete.

Prompts reflected each company's category, core use cases, and likely citation queries. Results were recorded by prompt, engine, company, and cited domain.

The testing was used to evaluate observed citation patterns and structural routing behavior. Because AI retrieval systems change continuously, these results represent patterns observed during the testing window, not a permanent citation state.

Data Sources and Quality Controls

Data sources

Appendix A

Research Methodology

Sources:

- **Ahrefs:** Organic traffic, traffic value, referring domains, rankings, and Top Pages
- **Google Search Console:** Branded/non-branded and CTR analysis for the limited cases where first-party access was available
- **G2 and supplementary directories:** candidate discovery only
- **Crunchbase:** Founding, funding, ownership, and M&A context
- **Scrunch:** AI citation testing across six engines

Interpreting Ahrefs estimates

Ahrefs traffic figures are modeled estimates rather than first-party session data. The study therefore emphasizes directional movement, relative change, position distribution, and cross-company differences rather than treating estimated traffic as an exact session count. Larger differences, repeated patterns across companies, and findings supported by multiple Ahrefs reports provide stronger evidence than isolated changes in a single estimated metric.

Quality controls

The following controls were applied across the research process:

- Inclusion and exclusion decisions were documented before analysis began.
- Classification thresholds were locked before the cohort distribution was known.
- The same formulas were applied uniformly across all 118 companies.
- Every manual classification override was recorded.
- All 15 deep dives used the same 14-field framework without selective abbreviation.
- Every deep dive required evaluation of competing explanations before mechanism assignment.
- Algorithm update and AI search expansion timelines were compiled before synthesis.
- Cross-company conclusions were developed only after all individual analyses and comparison briefs were complete.
- Low-scale, technically distorted, or unclassifiable cases were prevented from carrying full interpretive weight in the synthesis.

Scope and Interpretation

Observational evidence

The study identifies consistent associations, not experimental causality. Its conclusions are strongest where the same pattern appeared across multiple companies, categories, and analytical views.

Modeled search data
Ahrefs estimates are best used to compare direction, composition, and relative magnitude. Individual traffic figures should not be interpreted as exact first-party session counts.

Defined market sample

The cohort consists of B2B SaaS companies discoverable through major software directories and with sufficient organic presence for Ahrefs to model. The findings are most applicable to comparable B2B software companies and may not extend directly to very early-stage businesses or sites with little measurable search presence.

Defined study period

The research reflects the search environment from January 2025 through June 2026. Specific AI features and retrieval

Appendix B

Cohort Summary

This appendix summarizes the 118-company study cohort. For each company it reports category, traffic classification, traffic change, traffic trajectory, and Phase 3 analysis status.
Companies are grouped by category and sorted within each category by classification (Grew → Flat → Declined) and traffic change.

Phase 3 analysis key:

- ★ Deep dive — full 14-field mechanism analysis
- ◆ Comparison brief — structured paired comparison

Two companies (Semaphore and Kertos) began the study window with near-zero traffic, making percentage change mathematically unstable. Both are classified Grew based on the magnitude and direction of traffic acquired during the study window and therefore appear without a reported percentage. Semaphore remains a full-weight classifiable grower because page-level evidence supported confident assignment, while Kertos was excluded from the content growth-pattern distribution because its traffic base remained too small for defensible attribution.

Cohort totals 118 companies · 10 categories · 18-month window (Jan 2025 – Jun 2026) Grew — 44 companies (37%) Flat — 19 companies (16%) Declined — 55 companies (47%)

Customer Success · n = 13

Velaris * — Grew — +401.2% — Growth then Pullback Gainsight — Grew — +253.1% — Continuous Acceleration — ★ ◆ Akita — Grew — +185.8% — Continuous Acceleration — ★ ◆ Higher Logic — Grew — +32.2% — Continuous Acceleration ChurnZero — Grew — +26.1% — Growth then Pullback Rocketlane — Flat — +8.4% — Decline then Recovery Planhat — Flat — +6.3% — Decline then Recovery ClientSuccess — Flat — −3.8% — Decline then Recovery EverAfter — Flat — −6.4% — Decline then Recovery Custify — Declined — −34.6% — Growth then Pullback Vitally — Declined — −36.6% — Growth then Pullback Kapta — Declined — −39.6% — Early Peak then Decline UserGuiding — Declined — −80.4% — Growth then Pullback

Low-weight in content growth-pattern distribution (near-zero traffic base).

Sales Engagement/Sales Intelligence · n=11

ZoomInfo — Grew — +154.7% — Continuous Acceleration — ◆ Amplemarket — Grew — +92.2% — Continuous Acceleration Apollo.io — Grew — +35.5% — Continuous Acceleration Instantly.ai — Grew — +29.3% — Continuous Acceleration GetAccept — Grew — +15.7% — Growth then Pullback Salesloft — Flat — +7.1% — Growth then Pullback Mixmax — Declined — −24.2% — Early Peak then Decline Close — Declined — −26.5% — Early Peak then Decline Lemlist — Declined — −32.5% — Early Peak then Decline Cognism — Declined — −52.2% — Early Peak then Decline — ◆ Reply.io — Declined — −79.0% — Growth then Pullback

Marketing Automation · n = 11

Braze — Grew — +59.9% — Continuous Acceleration — ★ ◆ Customer.io — Grew — +51.0% — Continuous Acceleration Attentive — Grew — +42.3% — Continuous Acceleration Klaviyo — Flat — +8.9% — Volatile Brevo — Declined — −12.4% — Early Peak then Decline ContactPigeon — Declined — −14.5% — Volatile ActiveCampaign — Declined — −16.1% — Early Peak then Decline — ◆ CleverTap — Declined — −24.0% — Growth then Pullback Iterable — Declined — −45.1% — Early Peak then Decline GetResponse — Declined —

Appendix B

Cohort Summary

–60.9% — Early Peak then Decline Omnisend — Declined — –63.3% — Early Peak then Decline

HR / People Operations · n = 11

Rippling — Grew — +28.0% — Growth then Pullback — ★ ◆ HiBob — Grew — +18.9% — Growth then Pullback Deel — Flat — +6.6% — Volatile Workday — Declined — –16.2% — Growth then Pullback Gusto — Declined — –19.3% — Early Peak then Decline Paylocity — Declined — –31.8% — Steady Decline — ◆ BambooHR — Declined — –34.4% — Steady Decline Justworks — Declined — –35.0% — Early Peak then Decline TriNet — Declined — –44.0% — Steady Decline Remote — Declined — –58.1% — Early Peak then Decline Keka — Declined — –77.9% — Steady Decline

Customer Support / Help Desk · n = 11

Sobot † — Grew — +394.6% — Growth then Pullback TeamSupport † — Grew — +298.2% — Early Peak then Plateau Front — Grew — +14.5% — Continuous Acceleration BoldDesk — Declined — –14.9% — Growth then Pullback LiveAgent — Declined — –20.4% — Volatile Sprinklr — Declined — –31.8% — Early Peak then Decline Gorgias — Declined — –32.0% — Growth then Pullback Freshworks — Declined — –40.0% — Steady Decline — ★ Help Scout — Declined — –79.7% — Early Peak then Decline Hiver — Declined — –83.4% — Volatile EngageBay — Declined — –88.7% — Growth then Pullback

†Excluded from content growth-pattern distribution.

Finance / Spend Management · n = 11

SpendHound * — Grew — +436.6% — Continuous Acceleration Ramp — Grew — +112.5% — Continuous Acceleration — ★ ◆ Navan — Grew — +36.0% — Steady Growth — ★ ◆ BILL — Grew — +25.0% — Continuous Acceleration Corpay Complete † — Grew — +24.3% — Growth then Pullback Extend — Flat — –5.6% — Volatile Qonto — Declined — –10.7% — Volatile Procurify — Declined — –18.5% — Volatile Pleo — Declined — –19.2% — Early Peak then Decline Moss — Declined — –42.3% — Early Peak then Decline Coast — Declined — –51.1% — Growth then Pullback

*Low-weight in content growth-pattern distribution (near-zero traffic base).
†Excluded from content growth-pattern distribution.

Project Management/Collaboration · n = 11

Celoxis — Grew — +186.0% — Continuous Acceleration Notion — Grew — +12.5% — Growth then Pullback Airtable — Flat — +6.3% — Volatile Basecamp — Flat — +4.2% — Flat monday Work Management — Flat — –4.1% — Decline then Recovery Asana — Declined — –18.6% — Early Peak then Decline Wrike — Declined — –19.7% — Early Peak then Decline Teamwork.com — Declined — –28.4% — Volatile Quickbase — Declined — –29.9% — Steady Decline Float.com — Declined — –43.0% — Early Peak then Decline Hive — Declined — –49.3% — Early Peak then Decline

Developer Tools / DevOps · n = 11

Harness — Grew — +105.8% — Continuous Acceleration — ★ ◆ Octopus Deploy — Grew — +76.0% — Continuous Acceleration CloudBees — Grew — +71.9% — Continuous Acceleration — ★ ◆ CircleCI — Grew — +37.3% — Decline then Recovery — ★ ◆ Semaphore — Grew — N/A (near-zero base) — Continuous Acceleration JFrog — Flat — +2.2% — Flat GitLab — Flat — –1.9% — Volatile Bitbucket — Flat — –3.7% — Flat Perforce — Declined — –22.1% — Early Peak then Decline Bitrise — Declined — –24.7% — Volatile GitHub — Declined — –72.0% — Early Peak then Decline — ★ ◆

Appendix B

Cohort Summary

Security / Compliance · n = 14

Scrut Automation — Grew — +177.1% — Growth then Pullback Oneleet — Grew — +63.9% — Continuous Acceleration Vanta — Grew — +60.3% — Continuous Acceleration — ★ ◆ Hyperproof — Grew — +26.8% — Decline then Recovery — ★ Drata — Grew — +24.1% — Growth then Pullback SAI360 — Grew — +16.5% — Growth then Pullback Kertos † — Grew — N/A (near-zero base) — Growth then Pullback Secfix — Flat — +0.5% — Decline then Recovery ISMS.online — Flat — −5.3% — Volatile OneTrust Tech Risk & Compliance — Declined — −16.0% — Early Peak then Decline LogicGate — Declined — −21.6% — Early Peak then Decline Secureframe — Declined — −35.2% — Growth then Pullback — ◆ Sprinto — Declined — −39.2% — Decline then Recovery — ◆ Apptega — Declined — −41.5% — Steady Decline
†Excluded from content growth-pattern distribution.

Product Analytics / Product Management · n = 14

PostHog — Grew — +131.4% — Continuous Acceleration — ★ ◆ Mixpanel — Grew — +84.8% — Continuous Acceleration Pendo — Grew — +64.1% — Continuous Acceleration LaunchDarkly — Grew — +43.2% — Continuous Acceleration Canny — Grew — +42.3% — Volatile Glassbox — Grew — +33.3% — Steady Growth LogRocket — Grew — +19.8% — Decline then Recovery — ★ Amplitude Analytics — Flat — +2.8% — Growth then Pullback Fullstory — Flat — +1.4% — Decline then Recovery Productboard — Flat — −9.6% — Early Peak then Decline Aha! — Declined — −42.4% — Steady Decline ProductPlan — Declined — −60.0% — Steady Decline Userpilot — Declined — −81.2% — Early Peak then Decline — ◆ Whatfix — Declined — −91.3% — Early Peak then Decline

Content Growth-Pattern Distribution

Distribution among the 38 full-weight classifiable growing companies. Four companies were excluded because technical or unit-of-analysis issues prevented defensible attribution; two additional companies were treated as low-weight because of near-zero traffic bases. See Appendix A for methodology.
DE — Definitional/Educational: 23 companies (61%) B — Brand/Product-Demand: 7 companies (18%) M — Mixed: 7 companies (18%) P — Programmatic: 1 company (3%) T — Tools/Free Resources: 0 (—)

Growth-pattern distribution by trajectory:

Continuous Acceleration (n = 24) DE 15 (63%) · B 5 (21%) · M 3 (13%) · P 1 (4%)
Growth then Pullback (n = 8) DE 5 (63%) · B 1 (13%) · M 2 (25%) · P 0
Decline then Recovery (n = 3) DE 2 (67%) · B 0 · M 1 (33%) · P 0
Steady Growth (n = 2) DE 0 · B 1 (50%) · M 1 (50%) · P 0
Volatile (n = 1) DE 1 · B 0 · M 0 · P 0
DE appeared across all five trajectory types, including the two recovery trajectories. Brand/Product Demand appeared almost entirely in Continuous Acceleration.

Phase 3 Analysis Summary
Deep dives — 15 Comparison briefs — 9
Deep dives: Gainsight, Akita, Braze, Rippling, Freshworks, Ramp, Navan, Harness, CloudBees, CircleCI, GitHub, Vanta, Hyperproof, PostHog, LogRocket

Appendix C

Metrics & Calculations

This appendix defines the formulas, inputs, and interpretation thresholds for every quantitative metric used in the report. Conceptual definitions appear in the Glossary; research methodology is described in Appendix A.
Metrics are grouped by function: Traffic Composition → Portfolio Health → Position & Ranking → Classification.

Traffic Composition Metrics

These metrics decompose the headline traffic number to reveal what the composite obscures (Ch 3).

1. **Branded vs. Non-Branded Split**

What it measures The share of a program's organic traffic — and organic traffic growth — attributable to navigational (brand) queries versus non-branded category queries. These two components have different drivers and different implications for content program health.

Formula — share of total traffic

Branded traffic % = Branded organic clicks / Total organic clicks × 100

Formula — share of traffic growth

Branded % of growth = (Branded clicks T2 – Branded clicks T1) / (Total clicks T2 – Total clicks T1) × 100

Where T1 = start of measurement period, T2 = end of measurement period.
Data inputs Google Search Console (query-level click data with brand filter). Ahrefs and Semrush both provide branded/non-branded segmentation in their organic traffic exports.

Thresholds

< 40% — Growth is primarily content-driven — non-branded compounding 40–60% — Mixed signal — run both branded and non-branded trend separately

60% — *Headline primarily reflects product momentum, not content program compounding*

Cohort example Gainsight: approximately 65% of apparent traffic growth during the study window was attributable to a single branded keyword. Headline showed a growing organic program; non-branded content program had limited compounding. Rippling showed the inverse: branded traffic grew with product adoption while the non-branded content program experienced concurrent churn.

Note The metric is diagnostic, not evaluative — brand demand growth is a legitimate signal that requires a different response than content-driven growth.

2. Traffic Value vs. Raw Traffic Divergence

What it measures The gap between the rate of change in raw organic traffic and the rate of change in estimated traffic value over the same period. When these move together, traffic composition is roughly stable. When they diverge, the composition is shifting — toward lower-value queries (traffic up, value flat or down) or retaining higher-value positions despite volume loss (traffic down, value holds).

Formula

Traffic value divergence (pp) = Raw traffic % change – Traffic value % change

Where:
Raw traffic % change = (Raw traffic T2 – Raw traffic T1) / Raw traffic T1 × 100
Traffic value % change = (Traffic value T2 – Traffic value T1) / Traffic value T1 × 100

- Result is expressed in percentage points (pp). A positive divergence means raw

Appendix C

Metrics & Calculations

traffic grew faster than value (composition shifted toward lower-value queries). A negative divergence means value held better than volume (authority retained despite traffic loss).

Data inputs Ahrefs or Semrush domain overview (monthly traffic and traffic value series for the study window).

Thresholds

< 10 pp (either direction) — Composition roughly stable; headline is reasonably reliable 10–20 pp — Composition shifting; worth monitoring but not conclusive

20 pp — Significant composition shift; requires investigation before drawing conclusions from raw traffic headline

Cohort example ZoomInfo: Raw traffic grew 154.7%; traffic value grew 14.7% over the same period — a 140-point divergence. The volume growth was real; the quality of what was driving it had changed substantially. GitHub showed the inverse pattern: raw traffic fell significantly while traffic value declined at roughly half the rate — authority retained in high-value positions while lower-value volume was lost.

Portfolio Health Metrics

These metrics assess the internal composition and dynamics of a content portfolio over a defined period (Ch 5–6).

3. Gross Gain

What it measures The total organic traffic gained across all pages in a portfolio that increased during a measurement period, summed before netting against losses. Gross gain is a component input, not a standalone diagnostic — it is meaningful only in relation to gross loss.

Formula

Gross gain = Σ (Traffic T2 − Traffic T1) for all pages where Traffic T2 > Traffic T1

Data inputs Ahrefs Top Pages export at T1 and T2 (merged on URL). Minimum 3-month windows at each end to avoid single-month noise.

4. Gross Loss

What it measures The total organic traffic lost across all pages in a portfolio that declined during the same period, expressed as a positive absolute value. The numerator in the cannibalization rate formula.

Formula

Gross loss = Σ |Traffic T2 − Traffic T1| for all pages where Traffic T2 < Traffic T1

Data inputs Same as gross gain — page-level organic traffic for T1 and T2, filtered to pages with negative change.

5. Cannibalization Rate

What it measures The degree to which a content portfolio's gains are offset by losses elsewhere — the ratio of what the program is losing to what it is gaining. A high cannibalization rate means the program is running a treadmill: new production replaces existing losses rather than building net-new organic traffic.

Formula

Cannibalization rate = Gross loss / Gross gain × 100

Data inputs Gross gain and gross loss (see above). A 6–12 month window is recommended — shorter windows may mask longer-term treadmill patterns.

Thresholds

- < 30% — Healthy compounding — gains

Appendix C

Metrics & Calculations

substantially exceed losses 30–60% — Caution zone — losses are meaningful but growth is still net positive

60% — *Treadmill — production is primarily replacing losses, not building new traffic*
100% — *Net-negative investment — gross losses exceed gross gains; the program is shrinking*

Cohort examples LogRocket: 75.6% — for every unit of traffic gained, 0.76 units lost elsewhere. Aha!: 286.6% — for every unit gained, 2.86 units lost. Both programs had been producing at significant volume. The production output was feeding replacement, not compounding.

6. Net Gain (Net Traffic Change)

What it measures The headline traffic change: the difference between gross gain and gross loss. This is the number most commonly reported as "organic traffic growth." Meaningful only when interpreted alongside gross gain, gross loss, and the cannibalization rate that connects them.

Formula

Net gain = Gross gain − Gross loss

A positive net gain with a high cannibalization rate indicates replacement-driven growth, not compounding; a negative net gain with a low cannibalization rate may indicate deprioritization rather than structural decline.

7. Single-Page Concentration

What it measures The share of a portfolio's gross gains attributable to a single page over a defined period. The primary measure of concentration risk. When one page accounts for a disproportionate share of gross gains, any disruption to that page — algorithm update, AI Overview interception, competitive displacement — becomes a program-level event rather than an isolated page-level loss.

Formula

Single-page concentration = Largest positive page gain / Gross gain × 100
Gross gain is used as the denominator to avoid distortion when net gain is small or negative.

Data inputs Page-level traffic at T1 and T2 (same as gross gain calculation above).

Threshold

< 20% — Distributed — no single page is load-bearing for the program's growth 20–35% — Elevated — monitor for dependency

35% — *High concentration risk — a single-page event can reverse program trajectory*

Cohort context In the growing companies with the cleanest organic growth signals during the study window, no single page exceeded approximately 20% of gross gains. Braze, Navan, Harness, Vanta, and Ramp all showed broad distributed growth across dozens of pages. Programs chasing hero assets showed the opposite risk profile.

Position & Ranking Metrics

These metrics assess the distribution and trend of a portfolio's organic rankings across position bands (Ch 5).

8. Position Distribution

What it measures The count of ranking keywords (or pages) within each position band at a given point in time. Tracking how this distribution shifts over time is the primary observable for position consolidation.

Calculation

Appendix C

Metrics & Calculations

Position band count = Number of ranking keywords in positions [1–10 | 11–20 | 21–50 | 51+]

Run for two time points (T1 and T2) and compare band totals directionally.
Data inputs Ahrefs or Semrush (keyword count by position band at two time points).

9. Position Consolidation Signal

What it measures Whether a portfolio's ranking distribution is shifting toward top positions (1–10) while mid-tier and long-tail positions (21–50, 51+) contract. Position consolidation is a leading health signal — it indicates the portfolio is concentrating its authority at the positions with the highest click-through value.

Formula

Consolidation = (Positions 1–10 T2 – Positions 1–10 T1) > 0

AND (*Positions 21–50 T2 – Positions 21–50 T1) < 0*

This is a directional pattern check, not a pass/fail ratio. Magnitude matters: a program adding 5 top-10 positions while losing 200 mid-tier positions is consolidating differently than one adding 50 top-10 positions while losing 30 mid-tier positions. Interpret direction first, then scale.

Interpretation

Top-10 grows, 21–50+ contracts — Positive consolidation — program strengthening where it matters Top-10 flat, 21–50+ contracts — Long-tail erosion without top-position gains — monitor for further loss Top-10 contracts, 21–50+ flat/grows — Weakening — authority signals degrading at the positions that matter most All bands contract simultaneously — Broad-based decline — check for algorithm update alignment or architectural damage

Cohort example Hyperproof showed position consolidation as a leading recovery signal: top GRC positions (1–10) held and subsequently strengthened while long-tail compliance content declined. The structural floor was visible in the position distribution data before the headline traffic number confirmed recovery.

Classification Metrics

These metrics classify companies and study windows for cohort-level analysis (Appendix A, Appendix B).
10. Traffic Trajectory Classification
What it measures The primary classification for each company in the cohort: whether organic traffic grew, remained flat, or declined over the 18-month study window.

Formula

% change = (Traffic at window end – Traffic at window start) / Traffic at window start × 100

Primary classification bands

Grew — Net positive over the study window (directional threshold — sub-classifications apply) Flat — Within a narrow band around zero (sub-classifications apply) Declined — Net negative over the study window (sub-classifications apply)

Exact thresholds and sub-classifications (magnitude bands, volatility flags, recovery patterns) are reported in Appendix B.
Note on window definition Traffic at window start and window end are measured as 3-month averages (not single-month snapshots) to smooth seasonal variation. Specific window dates are defined in Appendix A.

Appendix C

Metrics & Calculations

11. Study Window Trough

What it measures The lowest monthly organic traffic figure recorded by a company during the 18-month study window. Used to anchor recovery assessments — recovery is measured from the trough, not from the window start.

Calculation

Study window trough = Min(monthly traffic) across all months in the study window

Data inputs Ahrefs Site Explorer (monthly traffic export for the study window). The trough is the minimum monthly value.

Why this matters A company that experienced a significant early decline and then recovered may show a negative net change over the full window while having experienced a substantial recovery from its lowest point. Measuring only start-to-end obscures the recovery. The trough is the anchor for recovery assessment; the window start-to-end is the cohort classification anchor.

12. Lift from Trough %

What it measures The percentage increase in organic traffic from a company's study window trough to its traffic level at the end of the study window. Used to assess recovery magnitude for companies that declined and subsequently recovered. Also referred to as Peak Lift % in the research workbook — the two names refer to the same calculation.

Formula

Peak lift % = (Traffic at window end − Study window trough) / Study window trough × 100

Interpretation Most useful for companies that declined and then recovered. For steady growers, Lift from Trough approximates the overall growth rate; for recovery cases, it isolates recovery magnitude independent of the window start.

Cohort example Hyperproof showed a peak lift that exceeded the magnitude of its initial decline — recovering not just to its pre-decline baseline but to a new all-time high within the study window. This is what the "top-position floor" recovery archetype looks like in the data: the structural floor held, and when external pressure receded, the recovery exceeded the loss.

Data Requirements Summary

Traffic Composition

- Branded/non-branded split: Google Search Console (query-level data with brand filter)
- Traffic value divergence: Ahrefs or Semrush (monthly traffic + traffic value)

Portfolio Health

- Gross gain / gross loss / cannibalization rate: Ahrefs Top Pages export at two time points
- Single-page concentration: same as above

Position & Ranking

- Position distribution / consolidation: Ahrefs or Semrush (keyword counts by position band at two time points)

Classification

- Traffic trajectory / study window trough / lift from trough: Ahrefs Site Explorer (monthly traffic export)

Minimum viable data pull for a single-program diagnostic:

- Ahrefs domain export: monthly traffic + traffic value, 18 months
- Ahrefs Top Pages export: two snapshots 12 months apart (page-level traffic)
- Google Search Console: 12 months of query-level data with brand filter

Glossary of Terms

This glossary defines every coined, redefined, or precision-critical term used in this report. Definitions here are the source of truth. Where a term was introduced in a specific chapter, that chapter is noted for cross-reference.

Terms are arranged in conceptual groups: Core Frame → Query & SERP Behavior → Traffic Analysis → Portfolio Health → Decline Classification → Recovery → Content Classification → Research Cohort Terms. An alphabetical index follows.

Core Frame

Durable demand: Organic demand built on content that compounds over time: content requiring continued engagement to deliver value, which develops ongoing relationships with its audience and renews itself because the need that produces the demand is structurally recurring. Resistant to AI Overview interception, algorithm volatility, and upstream demand substitution. Distinguished from fragile traffic by what happens to audience behavior after the first visit. (Introduced: Ch 1; Ch 11)

Fragile traffic: Organic traffic built primarily on satisfiable queries — content whose value can be fully extracted in a single interaction without requiring continued engagement. Includes template content, role and definition lists, brief how-to guides for well-documented processes, and procedural tutorials for mature technologies. Structurally at risk from AI Overview interception and upstream AI tool demand substitution. (Introduced: Ch 1)

Content durability: The extent to which a content asset retains its ability to generate organic demand despite changes in search behavior, AI retrieval, algorithm updates, or competitive conditions. In this report, durability refers to the persistence of audience demand and discoverability over time — not short-term ranking performance. Distinct from durable demand, which describes the audience behavior that a durable asset generates. An asset can be durable without currently producing large traffic volume; fragile traffic can produce large volume without durability. (Introduced: Ch 11; applied throughout)

Diagnosis before prescription: The methodological principle underlying this report: identify the mechanism producing an organic outcome before recommending an investment response. A program cannot be accurately diagnosed from headline figures alone. (Stated: Ch 1; applied throughout)

Query & SERP Behavior

Query satisfiability: The property of a search query that determines AI Overview exposure: whether the user's underlying task can be completed within the SERP — in the Overview response itself — without requiring the user to click through to an external page. More precise and more predictive than "informational vs. commercial intent." High satisfiability correlates with AI Overview CTR collapse; low satisfiability correlates with resilience. (Introduced: Ch 5; formalized: Ch 9)

Same-position CTR collapse: The data signature of AI Overview interception: a page's ranking position holds or improves while traffic to that page falls 80–99% on the same queries during the same period. Distinguishes AI Overview interception from competitive displacement (which produces position loss) and algorithm update damage (which produces position loss concentrated in mid-tier rankings). The clearest diagnostic signal available in standard rank-tracking data. (Phenomenon described: Ch 1, Ch 4; formally introduced: Ch 9)

AI Overview interception: One of the five

Glossary of Terms

decline mechanisms. Occurs when AI Overview responses in the SERP fully satisfy the user's task, preventing the click that would have occurred on the same ranking page without an Overview. Characterized by same-position CTR collapse. Mechanism operates at the SERP level, not upstream; the search still occurs but the click does not. Specific to satisfiable queries — does not broadly suppress informational content. (Introduced: Ch 4)

Complete query routing: The more extreme tier of AI search displacement, beyond AI Overview interception. AI systems route certain query types to specialist tools — live uptime monitors, real-time data sources, integrated calculators — rather than generating any content-based answer. No SERP ranking is relevant; the search behavior itself is redirected before the SERP interaction. Quality-independent and structural. (Introduced: Ch 10)

Structural routing The condition in which AI systems consistently route to institutional authorities or specialist tools for a given query type regardless of content quality, depth, or optimization. A content-quality problem cannot address structural routing; the routing occurs at the retrieval architecture level. Distinct from cases where editorial content can earn AI citation through information gain. (Introduced: Ch 10)

AI tool demand substitution: One of the five decline mechanisms. Occurs furthest upstream: users generate answers, templates, or procedural outputs directly from AI tools before initiating a search, so the query that would have driven traffic never occurs. Produces declining search volume on affected queries — not just declining CTR or rankings — distinguishing it from AI Overview interception. Template content, procedural how-to guides, and developer tutorials for mature technologies are most exposed. (Introduced: Ch 4)

Traffic Analysis

Traffic decomposition: The process of separating a headline organic traffic number into its component signals to reveal what the composite obscures. Three decompositions consistently changed the interpretation of headline figures in this study: branded vs. non-branded split, traffic value vs. raw traffic comparison, and category baseline comparison. (Introduced: Ch 3)

Branded vs. non-branded split: The separation of organic traffic into navigational queries (users searching for the brand or product by name) and non-branded category queries (users searching for topics, problems, or comparisons before knowing the product). Brand traffic reflects product adoption; non-branded traffic reflects content program compounding. When brand traffic accounts for more than 40% of growth in a period, the headline primarily reflects product momentum rather than content program health. (Introduced: Ch 3)

Traffic value divergence: A divergence of 20+ percentage points between the raw traffic trend and the traffic value trend over the same period. Signals that the composition of traffic is shifting — toward lower-value queries (traffic up, value flat) or retaining higher-value queries despite volume loss (traffic down, value holds). ZoomInfo's 140-point gap (154.7% raw traffic growth vs. 14.7% traffic value growth) is the most extreme example in the cohort. (Introduced: Ch 3)

Category baseline: The organic traffic trend of direct category peers over the same measurement window, used as a reference point for interpreting a program's absolute performance. A 30% decline in a category where the median declined 25%

Glossary of Terms

indicates different program-specific risk than the same decline in a category where peers held flat or grew. Determines whether a change is program-specific or market-wide. (Introduced: Ch 3)

Portfolio Health

Concentration risk: The fragility that results when a disproportionate share of organic gains comes from a single page or cluster. Threshold: when a single page accounts for more than 20% of gross gains, the portfolio has meaningful concentration risk — any change to that page (algorithm update, AI Overview interception, competitive displacement) produces a program-level event rather than an isolated page-level loss. (Introduced: Ch 5)

Position consolidation: The pattern in which positions 1–10 grow in count while positions 21–50 and 51+ contract. A leading health signal for a content program, not a warning sign. Indicates the portfolio is concentrating rankings in positions with the greatest click-through value, rather than accumulating mid-tier and long-tail positions with higher algorithm and AI Overview exposure. (Introduced: Ch 5)

Cannibalization rate: Gross loss as a percentage of gross gain in a content portfolio over a defined period. Calculated as: total traffic lost across declining pages : total traffic gained across growing pages. Under 30%: the portfolio is producing more new traffic than it is losing — healthy compounding. Over 60%: the portfolio is on a treadmill — new production increasingly replaces existing losses rather than building net-new traffic. Over 100%: net-negative investment, where every unit of gross gain is accompanied by more than one unit of gross loss. (Introduced: Ch 6)

Gross gain: Total organic traffic gained across all pages that increased during a measurement period, summed before netting against losses. Used to calculate cannibalization rate. (Introduced: Ch 6)
Gross loss Total organic traffic lost across all pages that declined during a measurement period, summed before netting against gains. Used to calculate cannibalization rate. The numerator in the cannibalization rate formula. (Introduced: Ch 6)

Single-page concentration: A specific measure of concentration risk: the share of gross gains attributable to a single page over a defined window. When a single page exceeds approximately 20% of gross gains, concentration risk is elevated. No growing company in the clean-growth cohort (Braze, Navan, Harness, Vanta, Ramp) showed a single page exceeding this threshold. (Introduced: Ch 5, Ch 6)

Content treadmill: The condition in which high cannibalization rates cause a program's production output to replace existing losses rather than build net-new traffic. The program appears productive — pages are being published, gross gains are occurring — but the net effect is displacement of existing traffic rather than compounding. Visible only when gross gain and gross loss are separated; the headline net number may appear modestly positive while the underlying dynamic is replacement. (Introduced: Ch 6)

Brand masking: The condition in which branded traffic growth — driven by product adoption, PR activity, or navigational search increases — offsets or obscures declining non-branded content program performance in the headline number. Programs experiencing brand masking appear healthier than their content programs are when viewed from the headline alone. Resolved by the branded vs. non-branded decomposition. (Introduced: Ch 3; coined: Ch 6)

Glossary of Terms

Decline Classification

Structural decline: A traffic decline produced by mechanisms not addressable through content production, link building, or standard SEO investment. Three structural patterns appeared in this study: pillar content collapse, architectural damage, and upstream demand substitution. (Introduced: Ch 8)

Tactical decline: A traffic decline produced by mechanisms addressable through content investment, authority building, or targeted optimization. Algorithm update damage and competitive displacement are the most common tactical decline mechanisms. Tactical declines can follow the recovery archetypes in Ch 7 when a structural floor exists. Distinguished from structural decline by whether standard recovery responses match the actual constraint. (Introduced: Ch 8)

Algorithm update damage: One of the five decline mechanisms. A sudden, simultaneous loss of rankings concentrated in the positions most vulnerable to spam or quality signals — typically positions 11–20, 21–50, and 51+ — within a narrow time window aligned to a known update date. Produces the most identifiable temporal signature in the data. May be tactical (recoverable through authority signals and content quality improvements) depending on what survived the update. (Introduced: Ch 4)

Competitive displacement: One of the five decline mechanisms. Gradual loss of rankings to competing domains earning stronger authority signals, newer content, or more topically comprehensive coverage. Distinguished from algorithm update damage (temporal, mid-tier concentrated) and AI Overview interception (same-position CTR collapse) by producing position loss across a longer timeframe without a clear external event trigger. (Introduced: Ch 4)

Architectural damage: One of the five decline mechanisms. Ranking losses produced by on-page content changes — redesigns, platform migrations, CMS changes, or content overhauls — that altered the targeting signals on existing pages. Referring domains may be retained; rankings are not, because the page-query matching signals changed. The fix is architectural, not editorial: publishing new content elsewhere does not restore targeting signals that were altered during a redesign. (Introduced: Ch 4; exemplified in structural decline context: Ch 8)

Pillar content: Foundational pages that define a program's core claim to topical authority in organic search — the pages establishing why the site exists for a given topic. Not supporting or peripheral content, but the primary definitional pages for a program's core categories (e.g., "what is SOC 2," "what is GRC"). When pillar content loses rankings, the supporting content that depended on pillar authority for its own ranking basis loses its structural support. (Introduced: Ch 8)

Pillar content collapse: The structural decline pattern in which a program's pillar content pages lose their rankings. When the foundation that anchored the program's authority is gone, recovery investment applied to peripheral or supporting content cannot restore what was the anchor. Hyperproof's contrast with Secureframe and Sprinto is the clearest evidence: Hyperproof's pillar GRC rankings survived while peripheral content declined; Secureframe and Sprinto lost the equivalent of Hyperproof's pillar layer.

Recovery

Structural floor The collection of surviving assets — pages, rankings, domain authority, or cluster authority — that continued to

Glossary of Terms

anchor a program's organic footprint during a decline period, and from which recovery became possible. Not every declining company retains a structural floor; when pillar content collapses or architectural damage reaches the foundation, the precondition for the Ch 7 recovery archetypes may be absent. (Coined: Ch 7)

Recovery archetypes: The three recovery archetypes observed among recovering companies in the cohort. Each archetype requires a specific surviving structural floor and proceeds through a distinct mechanism. (Introduced: Ch 7)
Content cluster breakthrough — Recovery driven by a content cluster already in the site consolidating upward simultaneously across a related topic area, without new content production. Requires: surviving domain authority and at least one area of surviving topical relevance. Typical timeline: 6+ months. (Case: CircleCI)

- Top-position floor — Recovery driven by surviving top-position content (positions 1–3) re-earning click volume as the basis for a single high-value ranking breakthrough. Requires: surviving top-position rankings on queries with genuine ongoing search intent. Fastest archetype: 2–4 months. (Case: Hyperproof)
- Re-ranking breakthrough on existing pages — Recovery driven by a single existing page re-ranking substantially on a high-volume query without new content production or detectable new backlink acquisition. Requires: surviving domain authority and an existing page with accumulated equity underperforming its ranking potential. Timeline: 4–6 months. (Case: CloudBees)

Content Classification

Engagement-requiring content: Content whose primary value requires the reader to read, think, and apply — content that cannot be extracted and used without continued engagement. Includes definitional and conceptual content, comprehensive methodology guides, proprietary frameworks, and original research. This content type consistently showed the greatest resilience across all three decline mechanisms in this study. Low query satisfiability; AI Overview responses cannot replicate the value of continued engagement. (Adjectival form: Ch 9; formally introduced: Ch 11)

Satisfaction-enabling content: Content whose primary value can be extracted in a single interaction without requiring continued engagement — a template, a brief definition, a list, a step-by-step procedure for a well-documented task. Corresponds to high query satisfiability. Most exposed to AI Overview interception and upstream AI tool demand substitution. Not inherently low-quality; the exposure is structural, based on the interaction pattern the content produces, not the craft with which it was written. (Formally introduced: Ch 11)

Content durability tiers: A four-tier classification of content types by their structural durability against the forces active in this study period. Tiers are based on the interaction pattern required to deliver value, not the production investment or SEO craft applied. (Introduced: Ch 11)

Tier 1 — Durable/Definitional: Content whose value is inherently ongoing and engagement-requiring. Foundational definitions, conceptual frameworks, methodology documentation, original research. Resilient across all three decline mechanisms.

Tier 2 — Durable with conditions: Content durable when specific conditions are met — typically, that it covers proprietary methodology, original data, or a topic with

Glossary of Terms

structurally recurring learning demand. Durable at the program level when managed for depth and authority; fragile when treated as commodity production.

Tier 3 — Durable with structural defense or proprietary data: Content in categories with high satisfiability that becomes durable only through structural differentiation — proprietary datasets, comparative data not reproducible elsewhere, tooling that requires genuine interaction. Without the differentiating layer, this tier defaults to Tier 4 exposure.

Tier 4 — Fragile: Content whose value is satisfaction-enabling: templates, brief definitions, role lists, procedural how-to guides for well-documented processes. Structurally exposed to AI Overview interception and upstream demand substitution. May retain traffic when competition for the query is low, but carries structural fragility under AI expansion conditions.

Information gain: The degree to which a piece of content adds to what is already known or available on a topic — going beyond what existing sources cover, introducing new evidence, synthesizing competing findings into a unified model, or providing data not found elsewhere. (Introduced: Ch 10)

Structurally renewable demand: Demand generated by content serving a need that is structurally recurring — where the audience that requires the content keeps arriving because the underlying need regenerates. Characteristic of foundational educational and definitional content, where new practitioners continuously enter a field and require the same foundational knowledge. (Introduced: Ch 5)

Research Cohort Terms

Traffic trajectory: The standardized classification applied to each company in the cohort based on its organic traffic direction and magnitude over the 18-month study window. Three primary classifications: Grew (44 companies), Flat (19 companies), Declined (55 companies). Within each, companies are further assigned a trajectory shape label (Continuous Acceleration, Growth then Pullback, Early Peak then Decline, etc.) describing the traffic curve. Details in Appendix B.

Study window: The 18-month observation period used to assess traffic trajectories and mechanism attribution across the 118-company cohort. All mechanism assessments (algorithm update damage, AI Overview interception, etc.) and recovery classifications are grounded in data from this window. (Defined: Appendix A)

Deep dive: A company-level investigation applying the Phase 3 analysis framework to identify the mechanisms behind a company's organic trajectory. Fifteen companies received deep dives; their findings form the primary evidence base for the report. (Defined: Appendix A)

Comparison brief: A structured comparison of two or more companies with similar category exposure but different traffic trajectories, used to isolate the variables that differentiated outcomes. Comparison briefs reduce confounding by controlling for category, time window, and competitive context. Nine comparison briefs contributed to the report's synthesis. (Defined: Appendix A)

Alphabetical Index

Algorithm update damage → Decline Classification
AI Overview interception → Query & SERP Behavior
AI tool demand substitution → Query & SERP Behavior

Glossary of Terms

Architectural damage → Decline Classification
Brand masking → Portfolio Health
Branded vs. non-branded split → Traffic Analysis
Cannibalization rate → Portfolio Health
Category baseline → Traffic Analysis
Complete query routing → Query & SERP Behavior
Competitive displacement → Decline Classification
Concentration risk → Portfolio Health
Content cluster breakthrough (archetype) → Recovery
Content durability → Core Frame
Content durability tiers → Content Classification
Content treadmill → Portfolio Health
Comparison brief → Research Cohort Terms
Deep dive → Research Cohort Terms
Diagnosis before prescription → Core Frame
Durable demand → Core Frame
Engagement-requiring content → Content Classification
Fragile traffic → Core Frame
Gross gain → Portfolio Health
Gross loss → Portfolio Health
Information gain → Content Classification
Pillar content → Decline Classification
Pillar content collapse → Decline Classification
Position consolidation → Portfolio Health
Query satisfiability → Query & SERP Behavior
Recovery archetypes → Recovery
Re-ranking breakthrough (archetype) → Recovery
Same-position CTR collapse → Query & SERP Behavior
Satisfaction-enabling content → Content Classification
Single-page concentration → Portfolio Health
Structural decline → Decline Classification
Structural floor → Recovery
Structural routing → Query & SERP Behavior
Structurally renewable demand → Content Classification
Study window → Research Cohort Terms
Tactical decline → Decline Classification
Top-position floor (archetype) → Recovery
Traffic decomposition → Traffic Analysis
Traffic trajectory → Research Cohort Terms
Traffic value divergence → Traffic Analysis