

SCAN TO
READ ONLINE[← PYC](#) • RESEARCHRESEARCH • METHOD / 30 QUESTIONS • 16 MONTHS • 5 ACCOUNTS

An audit is *thirty questions*, not a scan.

The questions a topical-authority audit asks of Search Console — put to sixteen months of real data across five live accounts. Every number is measured, sanitised by identity, and reproducible on your own data in an afternoon.

By Phil Yarrow • 17 min read • updated 2026-07-16

798

clicking queries • the
widest network

93%

non-brand clicks • the
top account

78

striking-distance
queries

860

zero-click pages •
retrieval overhang

26.8%

best top-3 click-
through

In one minute

THE WHOLE ARGUMENT, UP FRONT

- 01 An audit is a set of questions, not a scan. A topical-authority audit reads Search Console for one thing: is the site covering a subject completely, credibly and for long enough that ranking it is the safe bet? These are the thirty questions that answer that – put to sixteen months of real data across five live accounts.
- 02 Breadth of the query network is the first tell. The account with the widest coverage earns clicks on 798 distinct queries; the narrowest, 67. A wide, non-brand, distributed query network is the fingerprint of genuine authority – and it cannot be faked with volume.
- 03 Position, click-satisfaction and page architecture each answer a different question. Where does the coverage sit in the results? Does the site earn the click once it is shown? Is authority carried by one page or spread across the whole map? Two accounts rank well but lean on their own brand name; one is shown for everything and chosen for almost nothing.
- 04 The same questions expose the AI-era decoupling – and its exception. One account has 860 pages that get impressions and no clicks; another, whose position climbed most, saw its click-through rise while the era compressed everyone else's. The diagnosis is in the data, not the dashboard headline.
- 05 Every number here is real, sanitised by identity, and reproducible. This is the diagnostic discipline behind the field study and the guide – the questions we ask before we touch a line of content.

The data & method

The best SEO audit is not a tool that scores a page red or green. It is a fixed set of questions, asked in the same order every time, of the one dataset that records what a search engine actually did: **Search Console**. Get the questions right and the answers write the strategy for you.

The answers here come from sixteen months of Search Console performance exports across five live accounts — the same five in the [field study](#) — spanning e-commerce, local trades and B2B software. Every figure is measured from the raw query, page, device and search-appearance tables. Identities are NDA-protected, so each account is labelled by sector only, exactly as in the study.

One honest constraint shapes what can be claimed. Google caps the query and page exports at 1,000 rows each. So counts of *clicking* queries — all comfortably under a thousand — are true and uncapped; anything measured over the impression set is measured over each site's *top ~1,000* queries, and is described that way. The full [limitations](#) are listed at the end, not buried.

THE ONE QUESTION UNDERNEATH THE THIRTY

Is this site covering a subject completely, credibly, and for long enough that a search engine ranking it is the safe bet? Everything below is a way of asking that from a different angle.

SECTION I

The query network

How wide, how commercial and how distributed is the demand the site actually captures. Breadth of the query network is the first and hardest-to-fake signal of authority.

Query network · queries earning clicks

AUTO LOCKSMITH 798 QUERIES




SIGNAGE E-COMMERCE 536



GLAZING REPAIR 196



COMPLIANCE SAAS 80



PRESTIGE AUTOMOTIVE 67



Breadth of coverage. Uncapped — every count is under the 1,000-row export ceiling.

Non-brand share of clicks

SIGNAGE E-COMMERCE 93%



AUTO LOCKSMITH 76%



PRESTIGE AUTOMOTIVE 40%



GLAZING REPAIR 28%



COMPLIANCE SAAS 5%



Demand capture versus people typing the brand name.

01

How wide is the query network — how many queries earn clicks, not just impressions?

The clearest breadth signal, and unfakeable: Auto locksmith earns clicks on **798 distinct queries**, Signage e-commerce 536, Glazing repair 196, Compliance SaaS 80, Prestige automotive 67. A wide clicking network is coverage doing its job; a narrow one is a few pages carrying the site.

02

What share of clicks is non-brand demand versus people typing the brand name?

Signage e-commerce is **93% non-brand** — almost pure demand capture. Auto locksmith 76%, Prestige automotive 40%, Glazing repair 28%, Compliance SaaS just 5%. Non-brand is the traffic authority earns and the traffic that grows a business; brand clicks were coming anyway.

03

Is authority distributed across the network, or concentrated in a handful of head terms?

The top ten queries carry 44% of clicks at Auto locksmith and 53% at Signage e-commerce (distributed = real coverage), rising to **96% at Compliance SaaS** — where “the network” is essentially one query, the brand. Distribution is the

difference between a covered subject and a lucky page.

04

Does the demand the site captures overlap what the business actually sells?

At the demand-captured accounts the non-brand clicks are category and buyer-intent terms – the things that convert. At Compliance SaaS the non-brand demand exists only as impressions: **shown for the category, chosen for the name**. Coverage that does not capture commercial demand is a library, not a funnel.

SECTION II

Where the coverage sits

Coverage is worth what its position is worth. Where in the results does the mapped subject actually rank, and what is one push from breaking through?

Position distribution • top ~1,000 queries per account

ACCOUNT	TOP 3	TOP 10	PAGE 2+
Prestige automotive	20%	50%	50%
Auto locksmith	18%	44%	56%
Signage e-commerce	9%	55%	45%
Glazing repair	4%	16%	84%
Compliance SaaS	4%	15%	85%

Where the mapped subject ranks. Two accounts still sit mostly beyond page two – coverage in place, position climbing.

Striking-distance queries · positions ~8-20, real impressions

SIGNAGE E-COMMERCE

78 QUERIES

AUTO LOCKSMITH

31

COMPLIANCE SAAS

29

GLAZING REPAIR

8

PRESTIGE AUTOMOTIVE

2

The near-authority pipeline — the fastest return in the whole audit.

05

What share of the query set ranks top-3, top-10, or beyond page two?

Across each site's top ~1,000 queries: Prestige automotive is strongest (**20% top-3, 50% top-10**); Auto locksmith 18% / 44%; Signage 9% / 55%. Glazing repair and Compliance SaaS still sit mostly beyond page two (84-85% past position ten) — coverage in place, position still climbing.

06

How many “striking-distance” queries sit one push from page one?

Queries at positions ~8-20 with real impression volume — the near-authority pipeline: **Signage e-commerce has 78**, Auto locksmith 31, Compliance SaaS 29, Glazing repair 8, Prestige automotive 2. That is the fastest ROI on the whole audit: not new coverage, a nudge to coverage that already half-ranks.

07

Is average position improving over the sixteen months, or parked?

Every account measured is climbing: Glazing repair 46→24, Auto locksmith 38→19, Compliance SaaS 36→24, Prestige automotive from page two to the top of page one. Position is the leading indicator — the engine ranks you better before anyone clicks — and here it leads in the right direction across the board.

SECTION III

Earning the click

Ranking is being shown; the click is being chosen. At the same position, does the site earn the click — and where is it shown for things it never wins?

Click-through by position bucket • click satisfaction

ACCOUNT	POS 1-3	POS 4-10	POS 11-20
Glazing repair	26.79%	1.50%	0.37%
Prestige automotive	9.65%	1.09%	0.18%
Auto locksmith	7.09%	1.31%	0.33%
Compliance SaaS	6.67%	3.25%	0.04%
Signage e-commerce	2.61%	1.00%	0.15%

Same rank, different pull. Top-3 CTR ranges 10× across the five — visibility is not the same as being chosen.

08

At the same rank, does the site earn its clicks — is CTR at or above what the position should yield?

Top-3 click-through ranges from **26.8% at Glazing repair** (snippet and intent pulling hard) to just **2.61% at Signage e-commerce**, whose head terms are the most competitive and the most compressed by AI answers. Same rank, wildly different pull: visibility is not satisfaction.

09

Where are impressions high but click-through near zero — shown, but never chosen?

Compliance SaaS at positions 11-20 converts **0.04% of impressions** to clicks — effectively shown and ignored. That is coverage the engine is testing but users are not selecting: a signal to sharpen intent and snippet, not to publish more.

10

Do the SERP features the site owns actually earn the click?


Signage e-commerce wins both Product snippets and Merchant listings — but Merchant listings convert at **26.6% CTR against the rich snippet's 1.1%**. The high-intent shopping feature earns the click; the decorative one mostly does not. Owning a feature is not the same as being chosen through it.

SECTION IV

The page architecture

Authority spread across a mapped network is durable; authority balanced on one page is a single point of failure. How is the click-earning load distributed?

Authority concentration • share of clicks on the single top page



Lower is healthier — authority spread across the map, not balanced on one URL.

11

Is authority carried by one page, or spread across many?

The single top page accounts for just **19% of clicks at Prestige automotive** — healthy distribution across the map — versus 62% at Glazing repair and 66% at Compliance SaaS, where one URL is doing most of the work. Concentration is fragility: one algorithm shift on one page and the site moves.

12

How efficient is the coverage — how many pages earn clicks versus merely get impressions?

Compliance SaaS has 140 pages earning clicks and **860 shown-but-ignored**; Auto locksmith is lean at 68 earning against 39 ignored. A large ignored set is not always waste — it can be the outer section still ripening — but it is retrieval spent without return, and worth watching.

13

Are there pages with impressions and zero clicks that signal thin or mismatched coverage?

Thousands, portfolio-wide — most heavily at Compliance SaaS and Signage e-commerce. Read against the topical map, each is either outer-section coverage on its way up or a page answering the wrong intent. The audit's job is to tell those two apart, page by page.

SECTION V

Intent & source context

The same query means different things in different hands. Device, geography and query shape tell you the search context — and whether the coverage matches it.

Device split • clicks, and the intent it reveals

ACCOUNT	MOBILE	DESKTOP	READS AS
Auto locksmith	78%	19%	Urgent, in-the-field
Prestige automotive	77%	21%	Urgent, in-the-field
Signage e-commerce	60%	35%	Consumer shopping
Glazing repair	58%	38%	Home, mixed
Compliance SaaS	26%	72%	Considered B2B, at a desk

The channel mix confirms the intent the coverage has to serve.

14

What does the device split reveal about the search context?

Auto locksmith is **78% mobile** (locked out, on the move) and Prestige automotive 77% mobile — urgent, in-the-field intent. Compliance SaaS is **72% desktop**: a considered B2B search done at a desk. The channel mix confirms the intent the content has to serve.

15

Is the demand local, national or broad?

Signage e-commerce is **93% UK** — a national store — while the trades concentrate in their service areas. Geography tells you whether the topical map needs place-nodes (town-level service pages) or category depth. Get that wrong and you build the right coverage for the wrong map.

16

What is the dominant query intent, and does the coverage match it?


Transactional at Signage e-commerce and Auto locksmith, local-service at Prestige automotive and Glazing repair, informational-then-commercial at Compliance SaaS. Authority is only useful where the coverage answers the intent the queries actually carry — **the map has to be drawn to the demand, not to the keyword list.**

SECTION VI

Momentum & historical data

Authority is coverage plus track record. Is the network growing, is visibility converting to clicks, and did the coverage survive the updates that reset thin tactics?

The decoupling · Compliance SaaS – impressions vs clicks, indexed



Copper = impressions, grey = clicks. The widening gap is visibility the account is shown for but not chosen through – the AI-era decoupling, in one line.

17

Is the query network expanding over time, or static and decaying?

The single question this export cannot answer directly – GSC does not give query-by-month in a bulk pull – and the honest next data request. What the sixteen-month chart does show is total clicks and impressions climbing at four of five accounts, which only happens when new queries keep entering the network.

18

Is visibility converting to clicks, or decoupling from them?

Only one account decouples cleanly: Compliance SaaS, **impressions up ~230% while clicks rose ~18%** – the AI-era “great decoupling”. Glazing repair does the opposite, converting rising visibility into a rising click-through. Same era, opposite trajectories, and the difference is authority.

19

Is there a seasonal cycle the content calendar should anticipate?

The monthly series shows demand cycles – clearest in the seasonal trades – that a reactive calendar always publishes too late for. The audit turns that from a surprise into a schedule: **cover the seasonal outer-section before the season, not during it.**

20

Did the coverage hold through the 2025–2026 core updates, or reset?

It held and climbed. None of the five reset; the accounts built on comprehensive coverage rode the updates that were designed to reward exactly that. That is the whole thesis of authority in one observation — thin rankings evaporate on update day; covered ones gain.

SECTION VII

Cost of retrieval

Authority the engine cannot cheaply crawl and read does not count. How much indexed coverage is paying its way, and how much is retrieval spent for nothing?

21

How much indexed coverage never pays back the cost of crawling it?

Compliance SaaS carries **860 pages that earn impressions but no clicks** — the sharpest retrieval overhang in the set. Every one is crawled, rendered and indexed; a large fraction returns nothing. Lowering that ratio — prune, consolidate or fix intent — is cost-of-retrieval work that frees budget for pages that rank.

22

Is the click-earning page set concentrated enough to crawl efficiently?

Auto locksmith is efficient — a tight click-earning set with little dead weight — while Signage e-commerce (338 earning, 662 ignored) spreads crawl across a much larger surface. A big surface is fine if it is a real map ripening; it is a problem if it is duplication the engine keeps re-reading for nothing.

SECTION VIII

The strategic reads

The point of thirty questions is one paragraph of verdict per account – what the data says to do next, and what it says about the method itself.

The portfolio scorecard • the whole read in one matrix

ACCOUNT	NETWORK	NON-BRAND	POSITION	CTR	SPREAD	MOMENTUM	THE READ
Auto locksmith	●●●	●●●	●●●	●●●	●●●	●●●	Textbook authority – wide, non-brand, distributed, still climbing.
Signage e-commerce	●●●	●●●	●●●	●●●	●●●	●●●	Demand-captured and converting; head-term CTR is ceiling.
Prestige automotive	●●●	●●●	●●●	●●●	●●●	●●●	Position won – shift from building coverage to converting it.
Glazing repair	●●●	●●●	●●●	●●●	●●●	●●●	The compounding curve – position is the bottleneck, no coverage.
Compliance SaaS	●●●	●●●	●●●	●●●	●●●	●●●	Most AI-exposed – fix position and intent, not more pages.

●●● strong • ●●○ mixed • ●○○ weak, scored from the data above. Signals: query-network breadth, non-brand share, position, top-3 CTR, click spread across pages, momentum.

23

Which account shows the strongest topical-authority signal?

Auto locksmith: 798 clicking queries, 76% non-brand, distributed across the map (top ten only 44%), position climbing 38→19. Wide, commercial, distributed and improving – every fingerprint of authority at once. Its lead volume tells the same story from the business side.

24

Which is most exposed to the AI decoupling?

Compliance SaaS: 5% non-brand clicks, 860 zero-click pages, click-through collapsing to 0.04% deep in the results. It is shown for its category and chosen for its name. The fix is not more pages — it is position and intent on the coverage already indexed.

25

Where is the fastest return available?

Signage e-commerce's **78 striking-distance queries** — coverage already at positions 8–20 with real demand behind it. Nudging half-ranking pages onto page one returns faster than any new content, because the authority is already built; only the placement is missing.

26

Whose growth is brand-dependent versus demand-captured?

Glazing repair (28% non-brand) and Prestige automotive (40%) lean on their names; **Signage e-commerce (93%) and Auto locksmith (76%) are captured demand**. Brand-led growth is real but capped by brand awareness; demand-captured growth scales with coverage. The audit says which lever each account has left to pull.

27

Where is the problem “shown but not chosen” rather than “not shown”?

Signage e-commerce's compressed head terms and Compliance SaaS's deep-ranking category pages. Both are visibility without selection — **a snippet, title and intent problem, not a coverage problem**. Publishing more there would be answering a question nobody asked.

28

For each account, is position the bottleneck or is CTR?

Position is the bottleneck at Glazing repair and Compliance SaaS (the coverage is there, the rank is not). CTR is the bottleneck at Signage e-commerce's head terms (the rank is there, the click is not). **Naming which one, per account, is what turns an audit into a plan.**

29

Which account should shift from building coverage to converting it?

Prestige automotive: position won (top of page one, 50% of queries in the top ten), so the next unit of work is the click and the conversion, not another outer-section page. Knowing when to stop building is as valuable as knowing what to build.

30

Does the method predict the outcome?

Yes — and that is the whole point. Broad, distributed, non-brand, position-climbing coverage lands on exactly the accounts that produced the revenue and leads in the field study. The signal the audit reads and the business result it predicts are the same accounts. The questions work because the method does.

THE VERDICT

The signal the audit reads and the business result it predicts land on the same accounts. Broad, distributed, non-brand, position-climbing coverage is exactly what produced the revenue and the leads. The questions work because the method does — which is the only test of an audit worth running.

EVIDENCE • REPRODUCE IT

Run this on your own Search Console

None of these numbers need trusting on faith. Every one comes from a plain CSV export and a few lines of pandas — so you can run the same audit on your own site this afternoon. First, the exact definitions behind each figure, so

nothing rides on a fuzzy term:

Clicking query

A query with at least one click in the window. Breadth is counted this way because clicking-query totals sit under Google's 1,000-row export cap, so the count is true rather than truncated.

Non-brand

Any query whose text does not contain a brand token, matched case-insensitively. A close approximation — a few branded long-tail queries fall either side of the line.

Striking distance

Average position between 8 and 20 with at least 500 impressions: real demand, one page from page one. The fastest-return segment in the audit.

CTR by position

Clicks ÷ impressions within a position band (top-3 = positions 1–3). Reads click satisfaction independent of how many queries sit in the band.

Top-page concentration

Share of all clicks earned by the single highest-clicking URL. Lower is healthier — authority spread across the map, not balanced on one page.

Then the recipe. Export **Performance** → **Search results** → **Export** → **CSV** from Search Console (you get `Queries.csv` and `Pages.csv`), drop in your brand tokens, and run:


```

# pip install pandas

import pandas as pd

q = pd.read_csv("Queries.csv")          # Top queries, Clicks, Impressions, CTR, Position
BRAND = ["your brand", "yourbrand"]    # lowercase brand tokens
brand = q["Top queries"].str.lower().str.contains("|".join(BRAND))

# 1 · Query-network breadth — queries that earn clicks (uncapped; all under 1,000)
clicking = q["Clicks"].gt(0).sum()

# 2 · Non-brand share of clicks — the demand authority actually captures
nonbrand = q.loc[~brand, "Clicks"].sum() / q["Clicks"].sum()

# 3 · Striking distance — real demand, one push from page one
striking = (q["Position"].between(8, 20) & q["Impressions"].ge(500)).sum()

# 4 · Click satisfaction — CTR in the top three positions
top3 = q[q["Position"].between(1, 3)]
ctr_top3 = top3["Clicks"].sum() / top3["Impressions"].sum()

# 5 · Authority concentration — clicks on the single strongest page
p = pd.read_csv("Pages.csv")
concentration = p["Clicks"].max() / p["Clicks"].sum()

print("clicking queries :", clicking)
print("non-brand share  :", round(nonbrand * 100, 1), "%")
print("striking-distance:", striking)
print("CTR top-3       :", round(ctr_top3 * 100, 2), "%")
print("top-page share  :", round(concentration * 100), "%")

```

What this gives you — and what it doesn't

Five of the thirty questions, answered from your own data in one run — breadth, non-brand demand, the striking-distance pipeline, click satisfaction and authority concentration. It is descriptive, not causal: it reads the shape of your coverage, which is the part you control. The [field study](#) is where the causal argument, with before/after windows and significance tests, is made.

Limitations

The honest boundaries on what this data can and cannot claim. An audit that hides its limitations is a sales deck.

01 Five accounts, one operator, UK-centric. This is field evidence from a real portfolio, not a controlled study or a representative sample. It shows what the questions surface, not a population average.

02 Google caps the Search Console query and page exports at 1,000 rows each. Clicking-query counts (all well under 1,000) are therefore true and uncapped; anything measured over the impression set is measured over each site's top ~1,000 queries and is described that way throughout.

03 Brand versus non-brand is classified by matching the brand name in the query. It is a good approximation, not a perfect one – a handful of branded long-tail queries can fall on either side of the line.

04 Position and CTR are averages over the window, which hide within-period movement. The sixteen-month trend lines in the field study are the honest view of change over time; the tables here are the snapshot the audit reads.

05 Correlation, not proof of causation. These questions diagnose the shape of a site's coverage and the relationships in its data; they do not isolate a single cause. The field study is where the causal argument, with before/after windows and significance tests, is made.

Where to go next

This is the diagnostic discipline behind the method. The [field study](#) proves the thesis with sixteen months of before/after data; the [complete guide](#) teaches how to build the authority these questions measure. And the receipts are public:

The compounding curve – the field study

READ →

16 months, five businesses, real Search Console + GA4 data

Topical authority – the complete guide

READ →

What it is, why it beats keywords, and how to build it

Signage & personalisation e-commerce

£11,244

organic revenue · converts at 2× its traffic weight

Prestige-marque automotive specialist

5+ yrs

retained · 68% of traffic organic · #1 local for the marque

Domestic glazing & window repair

~5×

organic clicks in the recent quarter, YoY · page 5 → page 2

Auto locksmith & vehicle key replacement

505

qualified leads from organic · 77% of new customers

Compliance & quality-management software

Near-zero → p2

category head-term visibility built over 16 months



Thirty questions, answered with real data.

The diagnostic discipline behind the method — 30 questions put to five live accounts. To run it on your own Search Console, book a teardown at pyc.agency.

PYC • A TOPICAL-AUTHORITY AUDIT IN 30 QUESTIONS • VOL. 26 • JULY 2026
SCAN, OR VISIT [PYC.AGENCY/RESEARCH/TOPICAL-AUTHORITY-AUDIT](https://pyc.agency/research/topical-authority-audit) • [INFO@PHILYARROW.CO.UK](mailto:info@philyarrow.co.uk)