

Cadenta

cadenta.example · Cycle 3 · August 2026

68% AI
VISIBILITY

1·Foundations

2·Voice of Customer

3·Questions

4·Visibility

5·Action plan

6·Method & sources

CONTENTS

What's in this report

In the order the argument runs: who your buyers are, what they ask an AI, where you turn up in the answers, and what to change. Each section is measured against the one before it — so the fastest way to read this is straight through, once. Sections are numbered on every heading, so 1.4 is where the contents says it is.

1	Foundations		5
1.1	Who buys this	6	
1.2	Who you're compared against		6
1.3	What buyers evaluate	7	
1.4	Why buyers search		8
1.5	AI crawler access	9	
1.6	Content quality signals		9
1.7	Findings to fix	10	
2	Voice of Customer		13
3	Questions		15
3.1	Every question, by buyer	15	
3.2	Engineering manager · 30		
3.3	Operations lead · 30		
3.4	Agency owner · 30		
3.5	IT and security reviewer · 30		
3.6	Finance approver · 30		
4	Visibility		22
4.1	What this says	22	
4.2	The core insight		23
4.3	The tension	23	
4.4	Executive summary		23
4.5	What the visibility numbers say	24	
4.6	By buying stage		24
4.7	By buyer	24	
4.8	By pain point		24
4.9	By feature	25	
4.10	What the competitive numbers say		25
4.11	Share of voice	25	
4.12	Competitors you may not be tracking		26
4.13	Head to head — 84–81	26	
4.14	By engine		26
4.15	What the citation numbers say	26	
4.16	Where AI looks for answers		27
4.17	Who else gets cited	27	
4.18	Sources citing rivals but not you		27
4.19	Where you're absent	28	

5	Action plan	29
5.1	The plan	29
5.2	L1 — Fix the foundation	3
5.3	L2 — Optimize what exists	2
5.4	L3 — Build what's missing	1
6	Method & sources	34
6.1	How your foundation was built	34

SUMMARY

What we found

3 of your 6 findings are high severity. Those are the ones that change whether an AI engine can read and cite you at all; the rest are worth doing after them.

FINDINGS	WEAKEST SIGNAL	CRAWLER ACCESS	QUESTIONS
6 3 high · 2 medium · 1 low	0.39 Content freshness	4 of 4 AI answer-engine crawlers can reach you	150 across 5 buyer sets

FIX THESE FIRST

01 Pricing page renders its tiers client-side, so crawlers see an empty table

HIGH ENGINEERING 1-3 DAYS

02 Documentation sits behind a login wall

HIGH ENGINEERING 1 WEEK

03 The product page is published at two URLs with no canonical

HIGH ENGINEERING < 1 DAY

3 further findings — 2 medium · 1 low — with the same detail on each. [all 6 findings → · page 10](#)

Content freshness. Most scored pages haven't changed in months. Engines favour recent sources, so this is the cheapest lift you have. [all 4 signals → · page 9](#)

SECTION 1

1

Foundations

An engine recommends whoever best fits the buyer it thinks is asking. This is who we decided that buyer is, who you get compared to, and what they weigh — plus what a crawler can and can't read on cadenta.example today. Every number in this report is measured against these rows.

BUYERS 5 personas	COMPETITORS 6 in your category	FEATURES 8 buyers evaluate
PAGES 48 read of 93 found	FINDINGS 6 3 high · 2 medium · 1 low	

Why this section comes first. An AI engine can't recommend you until it can describe you, and it describes you using what it can read: your site, your competitors' pages, and the sources your category already trusts. So that is where we start, on purpose. Modelling your market from public evidence is fast and repeatable, and it puts a defensible map in front of you in days rather than the weeks a research project would burn.

Then we layer. What the engines already say about your category, and how your competitors actually position, are read into these rows before you see them. Your sales calls are the layer after that — the one that moves the ordering from what the category suggests your buyers weigh to what they were recorded saying, and the reason nothing below is attributed to a named buyer yet. The next section shows exactly what changes when they go in.

None of this is neutral description. Choosing which buyers count, which rivals you are measured against and which capabilities matter is choosing how you enter the arena — and every question in this crawl was written from these rows. Read it as the positioning call it is, and tell us where it misses.

1.1 Who buys this 5

The people who have to be convinced.

BUYERS

5

in the buying group we model

CAN BLOCK THE DEAL

3 of 5

hold a veto — convincing the others is not enough

Engineering manager

MANAGER

30 questions → · page 15

Engineering

Operations lead

DIRECTOR

CAN BLOCK THE DEAL

30 questions → · page 15

Operations

Agency owner

FOUNDER

30 questions → · page 15

Leadership

IT and security reviewer

SENIOR

CAN BLOCK THE DEAL

30 questions → · page 15

IT

Finance approver

CONTROLLER

CAN BLOCK THE DEAL

30 questions → · page 15

Finance

1.2 Who you're compared against 6

The vendors a buyer weighs you against, most directly compared first.

01

Kanbrix

CATEGORY LEADER

24% share of voice → · page 25

The default board tool, sold on how fast a team can start using it.

02

Sprintwell

DIRECT

21% share of voice → · page 25

Engineering-first planning with the deepest issue tracker in the category.

- 03

Northpath

ENTERPRISE

14% share of voice → · page 25

Consulting-led rollouts for regulated organisations that buy on governance.
- 04

Workstream

AUTOMATION

12% share of voice → · page 25

Workflow automation first, project views second.
- 05

Momentro

LIGHTWEIGHT

11% share of voice → · page 25

Minimal, opinionated, priced per workspace rather than per seat.
- 06

Pathgrid

ADJACENT

Resource and capacity planning expanding into delivery tracking.

1.3

What buyers evaluate

8

The capabilities buyers weigh when they compare you, ordered by how much of your question set each one earns.

CAPABILITIES	WHERE YOU'RE EXPOSED
8 2 strong · 3 moderate · 3 weak	3 of 8 rated weak or absent — highest-ranked is #1, Integration coverage

- 01

Integration coverage

WEAK

1 action → · page 29

Does it actually talk to the tools we already run?
- 02

Reporting and rollups

WEAK

Can I get a portfolio view without exporting to a spreadsheet?
- 03

Migration path

WEAK

1 action → · page 29

How much work is it to move four years of history across?
- 04

Permissions model

MODERATE

Can a client see their project and nothing else?

05 **Security and compliance** MODERATE 1 action → · page 29
Are you SOC 2, and will your DPA survive our legal review?

06 **Pricing transparency** STRONG 2 actions → · page 29
What does this cost at forty people, without a sales call?

07 **Automation rules** MODERATE
Can it move the ticket itself when the build passes?

08 **Onboarding effort** STRONG
Will the team still be using it in three months?

1.4 Why buyers search 5

The problems that put a buyer in the market, ordered by severity.

PAIN POINTS

5

2 high · 2 medium · 1 low

01 **Work is spread across four tools and no one view is true** HIGH
Status lives in three places and none of them agree.

02 **The last rollout was abandoned within a quarter** HIGH
We bought one of these already. Nobody opened it after March.

03 **Reporting to leadership is still assembled by hand** MEDIUM
I rebuild the same deck every Monday morning.

04 **Per-seat pricing punishes bringing in contractors and clients** MEDIUM
I'm not paying full price for someone who logs in twice a month.

05

Nobody can estimate the migration cost with a straight face

LOW

What happens to four years of tickets?

1.5 AI crawler access 4 / 4

Every crawler that feeds an AI answer engine can reach cadenta.example. CCBot is blocked, but it does not feed one — GPTBot, ClaudeBot, PerplexityBot, Google-Extended are each allowed by your robots.txt.

CRAWLER	STATUS
CCBot	blocked
GPTBot	allowed
ClaudeBot	allowed
PerplexityBot	allowed
Google-Extended	allowed

1.6 Content quality signals

AI platforms preferentially cite fresh, well-structured, deeply-covered pages they can extract clean passages from. Here's how the pages we read score.

Each bar is one score averaged across the 48 pages we analysed. The line at 0.75 is where we treat that average as healthy; below 0.50 we treat it as holding citations back. Individual pages can sit well below their own average — where a specific page or group is the problem, it is written up in the findings with its own number.

Content freshness 0.39

Most scored pages haven't changed in months. Engines favour recent sources, so this is the cheapest lift you have.

Content depth 0.71

On average your pages are substantial enough to quote, with room to go deeper on the commercial ones.

Heading structure 0.57

Heading structure is workable but inconsistent across your commercial pages.

Answer extractability 0.48

Answers are buried in prose an engine can't lift cleanly.

1.7 Findings to fix ⁶

Ordered by severity — high first. Each one is a concrete fix, with an owner and an effort estimate.

#	FINDING	SEVERITY	OWNER	EFFORT
01	Pricing page renders its tiers client-side, so crawlers see an empty table	HIGH	Engineering	1–3 days
02	Documentation sits behind a login wall	HIGH	Engineering	1 week
03	The product page is published at two URLs with no canonical	HIGH	Engineering	<1 day
04	Placeholder copy is live on two commercial pages	MEDIUM	Content	1–3 days
05	No integration page states what the integration actually does	MEDIUM	Content	1 week
06	Two changelog entries share a title tag	LOW	Content	<1 day

01 • Pricing page renders its tiers client-side, so crawlers see an empty table

HIGH

ENGINEERING

1–3 DAYS

Rendering

Owner · Engineering

Effort · 1–3 days

WHAT WE FOUND

The plan comparison is hydrated from a JSON call after load. In the raw HTML the table has headers and no rows.

WHY IT MATTERS

Price is the most-asked question in this category and the one buyers ask an assistant precisely because they do not want a sales call. Served as an empty table, the answer goes to whoever publishes a number.

RECOMMENDED FIX

Server-render the tier table and let the interactive toggle enhance markup that is already in the HTML.

Scope: The pricing page and the two comparison pages that embed the same component

02 • Documentation sits behind a login wall

HIGH

ENGINEERING

1 WEEK

Access

Owner · Engineering

Effort · 1 week

WHAT WE FOUND

All 340 documentation pages redirect anonymous requests to the sign-in screen, including the integration directory.

WHY IT MATTERS

Docs are what an engine quotes when a buyer asks whether a tool supports something specific. Behind a login they are invisible, and competitors' equivalents are not.

RECOMMENDED FIX

Open the reference and integration docs to anonymous readers and keep only account-specific pages gated.

03

● The product page is published at two URLs with no canonical

HIGH

ENGINEERING

< 1 DAY

Duplication

Owner · Engineering

Effort · < 1 day

WHAT WE FOUND

/product and /platform both return 200 with identical content and neither declares a canonical.

WHY IT MATTERS

Every citation and internal link splits across two addresses, so neither accumulates the authority one would.

RECOMMENDED FIX

301 /platform to /product and add a self-referencing canonical.

04

○ Placeholder copy is live on two commercial pages

MEDIUM

CONTENT

1-3 DAYS

Content quality

Owner · Content

Effort · 1-3 days

WHAT WE FOUND

A lorem-ipsum block and a heading reading “TEMPLATE — replace before launch” are live on the manufacturing and nonprofit landing pages.

WHY IT MATTERS

Placeholder text is a strong low-quality signal, and it sits on pages built to be cited.

RECOMMENDED FIX

Replace or unpublish both pages, then add a pre-publish check for placeholder strings.

05

○ No integration page states what the integration actually does

MEDIUM

CONTENT

1 WEEK

Depth

Owner · Content

Effort · 1 week

WHAT WE FOUND

The 62 integration pages carry a logo, a one-line blurb and a Connect button. None describe which objects sync, in which direction, or how often.

WHY IT MATTERS

“Does it sync two-way with our issue tracker” is a question engines answer from documentation. With nothing to lift, they answer it about somebody else.

RECOMMENDED FIX

Give each integration page a sync-direction table, an object list and a stated refresh interval.

06

○ Two changelog entries share a title tag

LOW

CONTENT

< 1 DAY

Metadata

Owner · Content

Effort · < 1 day

WHAT WE FOUND

Two 2025 posts both title as “Product update”.

WHY IT MATTERS

Duplicate titles make it ambiguous which page answers which question.

RECOMMENDED FIX

Retitle each to the change it actually describes.

SECTION 2

2

Voice of Customer

What your buyers say in their own words, validated against every row of section 1 — built from your recorded sales calls. This cycle ran on category research alone, so this page is a preview of what the recordings add.

When you share recorded calls, here is what this section holds:

What we heard

The patterns across your calls — each claim carrying the quotes that back it.

Who said it

The functions actually in the room, counted by accounts and people, set against the buyers your foundation models.

The pushback, in their words

Objections grouped by what was said, rather than by what we modelled.

What they ask you for

The requests that repeat across accounts.

What we'll optimise for

What the conversations change about the questions we run — row by row, each one linked back into your foundation.

SPECIMEN • AN INVENTED FINDING, DRAWN THE WAY YOUR REAL ONES ARE

Buyers hear “integration” and ask whether the sync is two-way

4 ACCOUNTS

The site says integrations; the calls say buyers only trust the ones that write back.

“If it can't push changes back, it's an import, not an integration.”

A buyer on one of your calls • 00:31:08 • linked to the recording

Only your buyers' and customers' own words count — never ours, never a partner's — and every quote links to the moment in the recording, so any claim can be checked against the conversation rather than taken on trust. Share the recordings and the next cycle builds this section from them, revalidates section 1 against it,

and re-weights the questions we run before anything is measured.

Questions

A buyer who already knows your name doesn't need an AI to find you. These are the questions asked by the ones who don't — a problem described, not a brand searched — and the engine picks who fits. Written from the foundation you approved, and the only questions your score is measured on.

QUESTIONS 150 from your foundation	BUYER SETS 5 grouped by persona	ENGINES 4 ChatGPT, Claude, Gemini, Perplexity	ANSWERS 600 captured & stored
CAPABILITIES 0 of 8 probed by these questions	PAIN POINTS 0 of 5 raised across the set	WIDEST SET 30 Engineering manager	

3.1 Every question, by buyer 150

These are the exact prompts, grouped by the buyer who asks them.

Engineering manager · Engineering · Manager

30 QUESTIONS

1	Best project management tool for a distributed engineering team?	action #4 →
2	Which PM tool has the best GitHub integration?	action #4 →
3	Project tracker that doesn't need a full-time admin?	
4	Cadenta vs Sprintwell for a 30-person engineering org?	
5	Best project management tool for a distributed engineering team? (21)	
6	Which PM tool has the best GitHub integration? (26)	
7	Project tracker that doesn't need a full-time admin? (31)	
8	Cadenta vs Sprintwell for a 30-person engineering org? (36)	
9	Best project management tool for a distributed engineering team? (41)	
10	Which PM tool has the best GitHub integration? (46)	
11	Project tracker that doesn't need a full-time admin? (51)	
12	Cadenta vs Sprintwell for a 30-person engineering org? (56)	
13	Best project management tool for a distributed engineering team? (61)	
14	Which PM tool has the best GitHub integration? (66)	
15	Project tracker that doesn't need a full-time admin? (71)	
16	Cadenta vs Sprintwell for a 30-person engineering org? (76)	
17	Best project management tool for a distributed engineering team? (81)	
18	Which PM tool has the best GitHub integration? (86)	
19	Project tracker that doesn't need a full-time admin? (91)	
20	Cadenta vs Sprintwell for a 30-person engineering org? (96)	
21	Best project management tool for a distributed engineering team? (101)	
22	Which PM tool has the best GitHub integration? (106)	
23	Project tracker that doesn't need a full-time admin? (111)	
24	Cadenta vs Sprintwell for a 30-person engineering org? (116)	
25	Best project management tool for a distributed engineering team? (121)	
26	Which PM tool has the best GitHub integration? (126)	
27	Project tracker that doesn't need a full-time admin? (131)	
28	Cadenta vs Sprintwell for a 30-person engineering org? (136)	
29	Best project management tool for a distributed engineering team? (141)	
30	Which PM tool has the best GitHub integration? (146)	

Operations lead · Operations · Director

30 QUESTIONS

1	How do you roll up status across twelve projects?	action #2 →
2	Project management tool with real portfolio reporting?	action #2 →
3	How long does it take to migrate off a legacy project tracker?	action #2 →
4	Which PM tool do operations teams actually stick with?	
5	How do you roll up status across twelve projects? (22)	
6	Project management tool with real portfolio reporting? (27)	
7	How long does it take to migrate off a legacy project tracker? (32)	
8	Which PM tool do operations teams actually stick with? (37)	
9	How do you roll up status across twelve projects? (42)	
10	Project management tool with real portfolio reporting? (47)	
11	How long does it take to migrate off a legacy project tracker? (52)	
12	Which PM tool do operations teams actually stick with? (57)	
13	How do you roll up status across twelve projects? (62)	
14	Project management tool with real portfolio reporting? (67)	
15	How long does it take to migrate off a legacy project tracker? (72)	
16	Which PM tool do operations teams actually stick with? (77)	
17	How do you roll up status across twelve projects? (82)	
18	Project management tool with real portfolio reporting? (87)	
19	How long does it take to migrate off a legacy project tracker? (92)	
20	Which PM tool do operations teams actually stick with? (97)	
21	How do you roll up status across twelve projects? (102)	
22	Project management tool with real portfolio reporting? (107)	
23	How long does it take to migrate off a legacy project tracker? (112)	
24	Which PM tool do operations teams actually stick with? (117)	
25	How do you roll up status across twelve projects? (122)	
26	Project management tool with real portfolio reporting? (127)	
27	How long does it take to migrate off a legacy project tracker? (132)	
28	Which PM tool do operations teams actually stick with? (137)	
29	How do you roll up status across twelve projects? (142)	
30	Project management tool with real portfolio reporting? (147)	

Agency owner · Leadership · Founder

30 QUESTIONS

1	Project management software that lets clients see only their project?	action #3 →
2	Best PM tool for an agency billing by retainer?	action #3 →
3	PM tool that doesn't charge full price for contractors?	action #3 →
4	Is Cadenta good for client work?	
5	Project management software that lets clients see only their project? (23)	
6	Best PM tool for an agency billing by retainer? (28)	
7	PM tool that doesn't charge full price for contractors? (33)	
8	Is Cadenta good for client work? (38)	
9	Project management software that lets clients see only their project? (43)	
10	Best PM tool for an agency billing by retainer? (48)	
11	PM tool that doesn't charge full price for contractors? (53)	
12	Is Cadenta good for client work? (58)	
13	Project management software that lets clients see only their project? (63)	
14	Best PM tool for an agency billing by retainer? (68)	
15	PM tool that doesn't charge full price for contractors? (73)	
16	Is Cadenta good for client work? (78)	
17	Project management software that lets clients see only their project? (83)	
18	Best PM tool for an agency billing by retainer? (88)	
19	PM tool that doesn't charge full price for contractors? (93)	
20	Is Cadenta good for client work? (98)	
21	Project management software that lets clients see only their project? (103)	
22	Best PM tool for an agency billing by retainer? (108)	
23	PM tool that doesn't charge full price for contractors? (113)	
24	Is Cadenta good for client work? (118)	
25	Project management software that lets clients see only their project? (123)	
26	Best PM tool for an agency billing by retainer? (128)	
27	PM tool that doesn't charge full price for contractors? (133)	
28	Is Cadenta good for client work? (138)	
29	Project management software that lets clients see only their project? (143)	
30	Best PM tool for an agency billing by retainer? (148)	

IT and security reviewer · IT · Senior

30 QUESTIONS

1	Is Cadenta SOC 2 compliant?	action #6 →
2	Project management tool with SAML and SCIM on every plan?	action #6 →
3	Which PM vendors will sign a DPA without an enterprise contract?	action #5 →
4	Where does Cadenta store customer data?	
5	Is Cadenta SOC 2 compliant? (24)	
6	Project management tool with SAML and SCIM on every plan? (29)	
7	Which PM vendors will sign a DPA without an enterprise contract? (34)	
8	Where does Cadenta store customer data? (39)	
9	Is Cadenta SOC 2 compliant? (44)	
10	Project management tool with SAML and SCIM on every plan? (49)	
11	Which PM vendors will sign a DPA without an enterprise contract? (54)	
12	Where does Cadenta store customer data? (59)	
13	Is Cadenta SOC 2 compliant? (64)	
14	Project management tool with SAML and SCIM on every plan? (69)	
15	Which PM vendors will sign a DPA without an enterprise contract? (74)	
16	Where does Cadenta store customer data? (79)	
17	Is Cadenta SOC 2 compliant? (84)	
18	Project management tool with SAML and SCIM on every plan? (89)	
19	Which PM vendors will sign a DPA without an enterprise contract? (94)	
20	Where does Cadenta store customer data? (99)	
21	Is Cadenta SOC 2 compliant? (104)	
22	Project management tool with SAML and SCIM on every plan? (109)	
23	Which PM vendors will sign a DPA without an enterprise contract? (114)	
24	Where does Cadenta store customer data? (119)	
25	Is Cadenta SOC 2 compliant? (124)	
26	Project management tool with SAML and SCIM on every plan? (129)	
27	Which PM vendors will sign a DPA without an enterprise contract? (134)	
28	Where does Cadenta store customer data? (139)	
29	Is Cadenta SOC 2 compliant? (144)	
30	Project management tool with SAML and SCIM on every plan? (149)	

Finance approver · Finance · Controller

30 QUESTIONS

1	What does Cadenta cost for 40 users?	action #1 →
2	Project management pricing compared without a sales call?	action #1 →
3	Cheapest project management tool that still does reporting?	action #1 →
4	Is per-seat or flat-rate PM pricing better value?	action #1 →
5	What does Cadenta cost for 40 users? (25)	
6	Project management pricing compared without a sales call? (30)	
7	Cheapest project management tool that still does reporting? (35)	
8	Is per-seat or flat-rate PM pricing better value? (40)	
9	What does Cadenta cost for 40 users? (45)	
10	Project management pricing compared without a sales call? (50)	
11	Cheapest project management tool that still does reporting? (55)	
12	Is per-seat or flat-rate PM pricing better value? (60)	
13	What does Cadenta cost for 40 users? (65)	
14	Project management pricing compared without a sales call? (70)	
15	Cheapest project management tool that still does reporting? (75)	
16	Is per-seat or flat-rate PM pricing better value? (80)	
17	What does Cadenta cost for 40 users? (85)	
18	Project management pricing compared without a sales call? (90)	
19	Cheapest project management tool that still does reporting? (95)	
20	Is per-seat or flat-rate PM pricing better value? (100)	
21	What does Cadenta cost for 40 users? (105)	
22	Project management pricing compared without a sales call? (110)	
23	Cheapest project management tool that still does reporting? (115)	
24	Is per-seat or flat-rate PM pricing better value? (120)	
25	What does Cadenta cost for 40 users? (125)	
26	Project management pricing compared without a sales call? (130)	
27	Cheapest project management tool that still does reporting? (135)	
28	Is per-seat or flat-rate PM pricing better value? (140)	
29	What does Cadenta cost for 40 users? (145)	
30	Project management pricing compared without a sales call? (150)	

150 questions · 5 buyers

EVERY ANSWER IS ON FILE

600 answers, across all 150 questions and all four engines.

Each question above was put to every engine, and each answer is stored word for word with the sources it cited, including the 48 questions you're absent from. They are in your account, not in this document — there are too many to print.

Read them at resonatelabs.co/account/. Anyone with access to the account can read them too; if a colleague should have access, tell us and we'll add them.

Visibility

When somebody asks an AI which vendor to use, it writes one answer and names a handful of companies. Being in that answer is the whole game, and it is not the same as ranking: there is no page two. Here is how often you make it, across 150 questions your buyers actually ask.

48 of those questions have no answer with you in it. They are listed one by one in your account, each with what every engine said instead.

Across 150 buyer questions, AI engines put you in the answer **68% of the time**, and you hold **18% share of voice** against Kanbrix's 24%.

AI VISIBILITY	WIN RATE	SHARE OF VOICE	INVISIBLE
68% 102 of 150	44% CI 39–49	#3 in category	48 questions you're absent from

4.1 What this says 3

- Cadenta appears in 68% of buyer answers but wins the recommendation in 44% of them. Visibility is not the bottleneck, preference is.
- Every capability question goes to a competitor. That single cluster accounts for 22 of the 48 questions Cadenta is absent from, and all of it is answered from documentation that sits behind a login.
- Four roundup pages supply a third of every answer in this category. Two omit Cadenta and two list last year's pricing.

4.2 The core insight

Known, but not quotable.

Engines can find Cadenta and describe it accurately. When a buyer asks something specific, about price, about whether an integration syncs both ways, about SOC 2, they reach for the vendor whose answer is sitting in readable HTML.

4.3 The tension

Strong when seen, absent when it counts.

The two numbers point in opposite directions and only one of them is a reach problem. Fixing resonance would be expensive and unnecessary. Fixing reach is a rendering and documentation exercise.

4.4 Executive summary

Cadenta is visible across two-thirds of the buyer questions in its category and holds third place in share of voice. The gap is structural rather than editorial: the pages that would answer a buyer's specific question are either rendered client-side, gated behind a login, or published without the detail an engine can lift. Answers that require a concrete fact therefore pick a competitor who publishes one.

The product-market fit signal is real. When Cadenta appears it wins 44% of contests outright. The problem is reach, not resonance.

At the early funnel, where buyers identify their problem and build a requirements list, Cadenta is invisible in 63% of queries. Buyers who never encounter it at discovery have no reason to shortlist it later. Two causes account for most of that.

First, the five editorial posts carrying informational coverage all date from 2025, outside the freshness window engines weight most heavily. Second, six weak or absent feature areas leave 60 queries with no addressable page on the site at all.


4.5 What the visibility numbers say

Visibility is not evenly distributed. It collapses at exactly the stages where a buyer decides who to consider.

Across 150 queries Cadenta appears in 68%, but that falls to 32% on problem-identification queries and 41% on requirements-building queries.


4.6 By buying stage ⁵

Buyers move through these stages in order. You are invisible in 62.8% of early-stage queries (94 of them) — the stage where consideration sets get formed.



4.7 By buyer ⁵

Each buyer type asks its own questions. A low number here means that person can't find you, whatever the overall score says.



4.8 By pain point ⁵

The problems buyers are trying to solve when they ask. Being absent on a pain is being absent from the reason they're searching.






* Too few queries in this segment for the number to be reliable.

4.9 By feature ⁶

The capabilities buyers weigh you on.




4.10 What the competitive numbers say

Cadenta beats the tools it is compared to and loses to the ones it is never compared with.


Head to head the record is 84–81, winning against Momentro and Pathgrid and losing to Kanbrix and Sprintwell, who hold the top two share-of-voice positions.

4.11 Share of voice ⁶



4.12 Competitors you may not be tracking 3

These brands show up in your buyers' answers without being on your competitor list.



4.13 Head to head — 84–81 5

COMPETITOR	WON	LOST	TIED	BOTH APPEAR IN
Momentro	24	8	2	32
Pathgrid	20	11	0	31
Werkstream	18	13	1	31
Sprintwell	12	22	0	34
Kanbrix	10	27	1	36

4.14 By engine 4

ENGINE	WON	SCORED	WIN RATE
ChatGPT	30	150	38%
Claude	43	150	52%
Gemini	38	150	45%
Perplexity	33	150	41%

4.15 What the citation numbers say

The engines already trust a small set of third-party sources in this category, and Cadenta is missing from most of them.

Three roundup and review sources supply 34% of all citations.

Cadenta's own domain ranks seventh with 67 citations across 14 pages.

4.16 Where AI looks for answers 7

Your own domain is cited 67 times, ranking #7 among all sources across 14 pages. These are the domains engines pulled from when answering your buyers' questions.

DOMAIN	CITATIONS
stackpicks.example	199
buildersforum.example	169
opsdigest.example	139
workflowwiki.example	123
kanbrix.example	93
sprintwell.example	83
cadenta.example	67

4.17 Who else gets cited 3

COMPETITOR	CITATIONS	CITED FROM
Kanbrix	93	stackpicks.example, buildersforum.example, opsdigest.example
Sprintwell	83	stackpicks.example, workflowwiki.example
Northpath	52	opsdigest.example, pmreview.example

4.18 Sources citing rivals but not you 5

Each of these is a specific, reachable placement: the engines already trust it, and it already covers your category.

SOURCE	CITES COMPETITORS
stackpicks.example	199
buildersforum.example	169
opsdigest.example	139
workflowwiki.example	123
pmreview.example	62

4.19 Where you're absent ⁴⁸

48 high-intent questions return an answer that never mentions you. 31 are won outright by a competitor, and 13 have no clear winner at all — those are the ones open to you.



SECTION 5

5

Action plan

Visibility is earned by being the easiest company in your category for an engine to quote — which is a content and structure problem, and therefore a solvable one. 6 moves, ranked, each tied to the questions it wins back and the pages it touches.

5.1 The plan 6

Ranked in the order we'd do them

Layer 1 moves the needle fastest; Layer 3 compounds.

L1 Fix the foundation **L2** Optimize what exists **L3** Build what's missing

1. Server-render the pricing table

L1

CRITICAL

ENGINEERING

WHY THIS MATTERS

Buyers ask an assistant what this costs precisely so they can avoid a sales call, and the answer they get names Momentro, whose price is in the HTML.

THE MOVES

- Server-render the tier table so the figures exist without JavaScript.
- Answer the 40-seat question in the first paragraph.
- Add an FAQ covering contractor seats and the annual discount.

WHO'S ASKING

Finance approver · Controller · Active Evaluation

PLATFORM PRIORITY

Chatgpt

Perplexity

Claude

WHY COMPETITORS WIN HERE

- Who Wins Today: Momentro, on nothing more sophisticated than publishing a number in server-rendered HTML.
- Why They Win: Their per-workspace price is a static string on the page, so an engine has something concrete to quote and a reason to quote it.

CLAIMS YOUR PAGE MUST MAKE

- State the price per seat per month at each tier, in the HTML.

- State what changes at the tier boundary, not only what it costs.
- Say plainly which capabilities require a sales conversation.

HOW TO FORMAT IT

Comparison Table

Direct Answer Paragraph

Faq

FORMAT NOTES

- Length Guidance: The answer in the first 60 words, then the table.

ON YOUR SITE

- Server-render the tier table and keep the monthly/annual toggle as progressive enhancement.
- Answer “what does this cost for 40 people” in the opening paragraph. Engines quote the opening, not the table.
- Add an FAQ block covering contractor seats, the annual discount and what the free tier excludes.

OFF-SITE

- Correct the outdated Cadenta pricing on the two roundup pages that still list the 2025 tiers.

NOTE

Publishing the number is the whole action. A page that says “contact us for pricing” cannot win a pricing query, and this category's buyers ask about price more than anything else.

QUESTIONS THIS ANSWERS

- What does Cadenta cost for 40 users?
- Project management pricing compared without a sales call?
- Cheapest project management tool that still does reporting?
- Is per-seat or flat-rate PM pricing better value?

FEATURES IT MOVES

Pricing transparency

PAGES IT TOUCHES

/pricing

2. Open the reference and integration documentation to anonymous readers

L1

HIGH

ENGINEERING

WHY THIS MATTERS

Your largest invisible cluster, 31 questions, is capability questions answered from docs. All 340 doc pages redirect to sign-in.

THE MOVES

- Serve the reference and integration directory to anonymous requests and gate only account-specific pages.
- Add a sync-direction table to each integration page so the capability answer is liftable.

QUESTIONS THIS ANSWERS

- How do you roll up status across twelve projects?
- Project management tool with real portfolio reporting?
- How long does it take to migrate off a legacy project tracker?

FEATURES IT MOVES

Integration coverage

PAGES IT TOUCHES

/docs

3. Publish a migration guide with a stated timeline

L2

HIGH

CONTENT

WHY THIS MATTERS

Operations buyers ask how long a move takes before they ask anything else, and Northpath owns that answer with a published 6-week plan.

QUESTIONS THIS ANSWERS

- Project management software that lets clients see only their project?
- PM tool that doesn't charge full price for contractors?
- Best PM tool for an agency billing by retainer?

FEATURES IT MOVES

Migration path

PAGES IT TOUCHES

/migrate

4. Get corrected into the category roundups

L2

MEDIUM

PR

WHY THIS MATTERS

Perplexity draws 40% of its answers in this category from four roundup pages. Two omit Cadenta and two list your 2025 pricing.

QUESTIONS THIS ANSWERS

- Best project management tool for a distributed engineering team?
- Which PM tool has the best GitHub integration?

FEATURES IT MOVES

Pricing transparency

5. Add FAQ schema to the security page

L1

MEDIUM

ENGINEERING

WHY THIS MATTERS

Your SOC 2 and data-residency answers sit in prose an engine has to infer structure from.

ROOT CAUSE

The page was written for a security questionnaire, not as an answerable document.

WHY THIS ITEM

Schema is the cheapest way to make an answer liftable. It changes nothing a human sees.

WHAT WE FOUND

The security page answers eleven compliance questions in body copy with no FAQPage markup.

WHY IT MATTERS

Engines lift structured answers far more readily than prose, and the IT reviewer persona is the one most likely to be researched by assistant rather than by call.

STEP BY STEP

- Add FAQPage JSON-LD with one entity per compliance question.
- Mirror the visible copy exactly. Schema that disagrees with the page is a trust penalty.
- Validate with the Rich Results test before publishing.

HOW TO VERIFY IT WORKED

- Rich Results test passes with 11 questions detected.
- Visible copy and schema answers match word for word.

WHAT THIS UNLOCKS

- Makes the security page quotable as a source for the enterprise comparison queries.

QUESTIONS THIS ANSWERS

- Which PM vendors will sign a DPA without an enterprise contract?

FEATURES IT MOVES

Security and compliance

PAGES IT TOUCHES

/security

6. Build a presence in the practitioner community threads

L3

LOW

PR

WHY THIS MATTERS

Community discussion is the second most-cited source type in this category and you appear in none of it.

QUESTIONS THIS ANSWERS

- Is Cadenta SOC 2 compliant?
- Project management tool with SAML and SCIM on every plan?

6 of 6

SECTION 6

6

Method & sources

How these numbers were produced, so they can be checked.

6.1 How your foundation was built

We build this from category research: your own site, the sites of the vendors buyers weigh you against, how the category is listed and described publicly, and what reviewers say on the public review platforms. You then review it and correct anything we got wrong, and nothing is measured until you approve it.

That makes it our reading of your market, checked by you — not a claim about what your buyers said, because on this cycle we have not heard from them directly. Recorded sales calls are how that changes. Here is what it looks like, on an invented row: the same line of section 1, before and after your buyers' own words go in.

THIS CYCLE · CATEGORY RESEARCH, REVIEWED BY YOU

Renewal arrives before anyone can prove it worked

HIGH

RESEARCH ONLY

The champion who bought it is left defending the spend with anecdotes.

NEXT CYCLE · THE SAME ROW WITH YOUR CALLS IN

Renewal arrives before anyone can prove it worked

HIGH

3 ACCOUNTS SAID THIS

PRIORITY #1

The champion who bought it is left defending the spend with anecdotes.

"We had no number to put in front of finance at renewal, so it got cut."

A buyer on one of your calls · 00:14:22 · linked to the recording

What your buyers confirm stays. What they contradict is changed. What they raise that the research missed is added — every quote carrying the speaker, the account and the second it was said. Share the recordings and section 1 is rebuilt against them on the next cycle, before anything is measured.

Sources we read directly:

- Your site — [cadenta.example](#) (48 of 93 pages read)
- Your sitemaps — the sitemap index and all 4 child sitemaps (131 URLs, 38 URLs excluded)
- A live crawler reachability check — 5 crawler user-agents requested your homepage directly, rather than reading your robots.txt alone

Each question is then put to 4 AI engines (ChatGPT, Claude, Gemini, Perplexity) and the full answer is recorded along with the sources it cited. That's 600 answers in this cycle, from 150 questions. Visibility is the share of questions where an engine puts you in the answer. Win rate is graded, not binary: a solo top pick scores 100%, a strong second 50%, a place in a list 25%, a passing mention 10%, and absence 0 — averaged across every question, including the ones you're absent from.

The technical findings come from a crawl of 48 of the 93 pages we could discover on your own site. Archive, tag and legal pages are excluded from that count, and because the site holds more pages than one crawl reads, the crawl always includes your most recently updated pages alongside every commercially relevant one, up to the crawl's page budget. Each page is checked for crawler access, freshness, structure and how cleanly a passage can be extracted and quoted.

Crawl	cadenta-c3
Data as of	July 26, 2026
Questions	150
Engines	ChatGPT, Claude, Gemini, Perplexity
Answers captured	600



This report was produced by Resonate Labs.

We measure how AI engines answer the questions buyers actually ask — across ChatGPT, Claude, Perplexity and Gemini — and hand back what to change, in the order it is worth changing.

resonatelabs.co

