Railway: AEO and AI Visibility Strategy

30-Day Campaign

Establish a reliable AEO measurement system, identify high-value opportunities, and ship improvements that make Railway more discoverable and useful. Build a repeatable process that connects content quality to growth.

Railway: AEO and AI Visibility Strategy **1**
- Phase 0 — Initial alignment and setup 2
  - Objectives 2
  - Initial requirements 2
  - Deliverable 2
- **Phase 1 — Research and baseline measurement** **3**
- Objective 3
  - 1.1 Prompt research 3
  - Initial query clusters 3
  - 1.2 AI platform coverage 4
  - 1.3 What to measure 4
  - 1.4 Segment the findings 5
  - 1.5 Technical accessibility baseline 5
- **Phase 2 — Identify opportunities** **6**
- Objective 6
- 2.1 Technical opportunities 6
  - Consideration: JSON-LD and llms.txt 7
- 2.2 Content opportunities 7
  - Content quality checklist 7
  - Content formats to prioritize 8
- 2.3 Competitive and market opportunities 8
  - Deliverable by Day 10 8
- **Phase 3 — Optimize and ship** **9**
- 3.1 Prioritization framework 9
  - Opportunity prioritization 9
- 3.2 Initial execution workstreams 9
  - Workstream A — Improve high-value existing pages 9
  - Workstream B — Create or improve technical examples 10
  - Workstream C — Technical site improvements 10
  - Workstream D — Editorial standards 11

Phase 4 — Measure expected impact **11**
Objective 11
4.1 Before-and-after comparison 11
Suggested evaluation table 12
Important measurement caveat 12
4.2 ROI and learning 12
Phase 5 — Iterate and establish the next roadmap **13**
Objective 13
Activities 13
Final deliverables 13

Phase 0 — Initial alignment and setup

Timing: Before or at the beginning of Week 1

Objectives

Understand Railway's product, customers, existing content ecosystem, and available measurement tools before making optimization decisions.

Initial requirements

- Develop a working understanding of Railway's product offering and cloud infrastructure.
- Become familiar with Railway's documentation, blog, templates, deployment workflows, and public-facing resources.
- Understand the distinction between:
    - Search engines and traditional web crawlers.
    - AI search/retrieval systems.
    - LLM chat interfaces and user-triggered retrieval.
    - Agentic workflows and how they may interact with public content.
- Understand the intersection of SEO and AEO.
- Learn the basic mechanics of crawling, rendering, indexing, retrieval, and citation.
- Gain access to Gauge and relevant analytics or technical monitoring tools.

Deliverable

A short internal orientation document containing:

1. Railway's primary audiences and use cases.
2. Core product capabilities and terminology.
3. Existing content and documentation inventory.

4. Available measurement tools and baseline access.
5. Initial hypotheses about where visibility opportunities may exist.

Do not assume that every AI system uses the same crawler, ranking process, or citation mechanism. The measurement framework needs to account for platform differences.

Phase 1 — Research and baseline measurement

Timing: Days 1–7

Objective

Establish a measurable baseline of Railway's visibility in AI-generated answers and identify which queries and user needs represent the greatest opportunity.

1.1 Prompt research

Use Gauge to monitor Railway's presence in AI-generated responses and develop an evaluation set of approximately 100–1,000 queries, depending on available time, query diversity, and evaluation cost.

Start with a smaller, carefully curated set, then expand it based on early findings rather than assuming that a large number of prompts automatically produces a better baseline.

Initial query clusters

Migration and alternatives

- How do I migrate from Heroku?
- Heroku alternatives for a small SaaS.
- How do I move an existing application to Railway?

Application deployment

- How do I deploy a Node.js application?
- How do I deploy a Docker container?
- How do I deploy Next.js?
- How do I deploy Python/FastAPI?
- How do I deploy Rails?

Infrastructure and services

- How do I deploy PostgreSQL?
- How do I deploy Redis?
- How do I deploy a monorepo?
- How do I set environment variables?
- How do I configure custom domains?

Operations and troubleshooting

- How do I configure health checks?
- How do I deploy to a specific region?
- How do I debug a failed Railway deployment?
- How do I scale a service?
- How do I set up preview environments?

Emerging workflows and comparison

- How do I deploy an MCP server?
- What is the best platform for deploying a Next.js API?
- How do I deploy FastAPI with PostgreSQL?
- How do I deploy a Docker container cheaply?

1.2 AI platform coverage

Run the evaluation set periodically across relevant platforms, such as:

- ChatGPT
- Gemini
- Perplexity
- Claude
- Microsoft Copilot
- Grok

Treat these as separate measurement environments rather than assuming that a result from one platform generalizes to every other system.

1.3 What to measure

Dimension	Measurement question
Mention visibility	Is Railway mentioned in the response?
Source visibility	Is a Railway URL cited or used as a source?
Recommendation	Is Railway recommended when the query requests a solution or platform?

URL coverage	Which Railway page is cited?
Capability accuracy	Is the referenced Railway feature described correctly?
Competitive presence	Which competitors appear in the answer?
Evidence	What sources or claims appear to support the answer?
User intent	What type of user need does the query represent?

Distinguish being mentioned from being recommended, because a brand can appear in an answer without being the preferred solution.

1.4 Segment the findings

Break down results by:

- Geography.
- Persona (individual developer, startup team, experienced engineer).
- AI platform/model.
- Language.
- Topic and use case.
- Search intent.
- Competitor set.

Look for patterns such as:

- Queries where Railway is not mentioned.
- Queries where Railway is mentioned but the relevant documentation is not cited.
- Queries where Railway's capabilities are misrepresented.
- Queries where competitors appear more consistently.
- Topics where Railway has strong documentation but weak answer visibility.

1.5 Technical accessibility baseline

Alongside prompt research, investigate whether relevant Railway pages are technically accessible and interpretable by the systems that matter.

Areas to examine:

- Crawler access and relevant robots rules.
- Server responses and access logs.
- Rendering and JavaScript dependencies.
- HTML structure and page content extraction.
- Canonical URLs, indexing directives, and sitemaps.

- Page performance and mobile rendering.
- Structured metadata and JSON-LD.

Deliverable by Day 7: A baseline report with the evaluation set, measurement methodology, initial findings, and a prioritized list of potential opportunities.

Phase 2 — Identify opportunities

Timing: Days 5–10

Objective

Translate the baseline findings into a prioritized list of content, technical, and product education opportunities.

The central question:

> Where is Railway losing the opportunity to be discovered, understood, cited, or recommended—and what is the smallest effective intervention?

Classify opportunities into three main areas.

2.1 Technical opportunities

Prioritize the following audit areas:

Crawler access and retrieval

- Review robots.txt and crawler policies.
- Distinguish AI search/retrieval bots from training-related crawlers.
- Inspect access logs for relevant crawlers where available.
- Check whether important content is blocked or inaccessible.

Content rendering and extraction

- Check whether important content is available in the initial HTML.
- Identify critical information hidden behind JavaScript.
- Test server-side rendering, content extraction, and page response behavior.
- Investigate broken links, redirects, and inconsistent content delivery.

Search infrastructure

- Review canonicalization and indexing directives.

- Check sitemap coverage and important URL discoverability.
- Audit metadata and relevant structured data.
- Review internal linking between documentation, guides, templates, and product pages.

Performance and mobile usability

- Investigate slow page loads and performance bottlenecks.
- Check mobile rendering and layout issues.
- Prioritize problems affecting access to useful content, rather than treating a generic speed target as a standalone AEO objective.

Consideration: JSON-LD and llms.txt

Include JSON-LD and a curated `llms.txt`. Do not assume that implementing either automatically improves AI rankings or citations.

- Structured data can help describe entities and content to search engines, but it is not a guaranteed AI citation lever.
- `llms.txt` is a proposed convention for organizing useful site information for language models. Its effectiveness as a direct ranking or citation mechanism remains uncertain.
- Robots directives can affect crawler access, but the policies and behavior of different bots are not identical.

2.2 Content opportunities

Assess the content itself based on whether it answers the user's question clearly, accurately, and usefully.

Content quality checklist

Area	What to evaluate
Relevance	Does the page address a specific user need?
Structure	Are the main questions and answers easy to identify?
Terminology	Does the page use the language developers actually use?
Answer clarity	Are important answers concise, explicit, and self-contained?
Depth	Does the content support the intended audience and complexity?
Accuracy	Are instructions and product capabilities correct?
Freshness	Is the content maintained as the product evolves?

Internal linking	Can users and crawlers reach related resources?
Authority	Does the page provide evidence, examples, or firsthand expertise where appropriate?
Accessibility	Are images, diagrams, and other content elements accessible?

Content formats to prioritize

Investigate opportunities across Railway's owned assets:

- Documentation and deployment guides.
- Blog posts.
- Tutorials and technical walkthroughs.
- Standalone application examples.
- Migration guides.
- Troubleshooting resources.
- Comparison or alternative pages, where useful and factually supportable.
- Feature and workflow documentation.

2.3 Competitive and market opportunities

Identify where competitors are appearing in answers that could reasonably include Railway.

For each competitive gap, investigate:

1. What question is the user asking?
2. Which competitor or competitors are appearing?
3. What evidence supports their inclusion?
4. Does Railway have a comparable capability?
5. Does Railway have a useful, accurate page covering that use case?
6. Is the gap caused by content, technical accessibility, product limitations, or another factor?

Deliverable by Day 10

An opportunity backlog with:

- Query or user intent.
- Relevant Railway page or missing content area.
- Observed problem.
- Proposed intervention.
- Expected impact hypothesis.

- Technical or editorial complexity.
- Dependencies and owner.
- Measurement plan.

Phase 3 — Optimize and ship

Timing: Days 10–23

3.1 Prioritization framework

Use a simple scoring model to organize decisions, without treating the score as a precise prediction of outcomes.

Opportunity prioritization

Prioritize changes where there is a clear user need, an observed visibility or content gap, and a feasible path to shipping and measurement.

- Potential impact: High / Medium / Low
- Relevance to strategic queries, user value, business importance, and the size of the identified gap.
- Effort and complexity
- Quick win / Project
- Engineering effort, editorial effort, dependencies, and validation requirements.
- Evidence strength
- Observed / Hypothesis

3.2 Initial execution workstreams

Workstream A — Improve high-value existing pages

Select a small number of pages from the baseline that are relevant to high-value queries and have identifiable weaknesses.

Potential improvements:

- Clarify the primary user question and answer.
- Refine headings and page structure.
- Add missing prerequisites and implementation steps.
- Improve examples and error handling.
- Add links to relevant documentation.
- Correct outdated or ambiguous product information.

- Improve page-level metadata and technical discoverability where justified.

Example: If research finds that developers frequently ask how to deploy FastAPI with PostgreSQL, examine the existing Railway resources for this workflow and identify whether the content sufficiently covers setup, deployment, environment configuration, connection details, and troubleshooting.

Workstream B — Create or improve technical examples

The role explicitly includes building standalone application examples. Use these to demonstrate actual Railway workflows and create content that is useful for both developers and answer engines.

Potential examples to evaluate:

- A deployable FastAPI + PostgreSQL application.
- A Next.js application with a backend service.
- A Docker-based deployment example.
- A monorepo deployment workflow.
- An MCP server deployment example.

The exact examples should be selected after reviewing user demand, existing assets, and the gaps discovered in Phase 1.

For each example, include:

- A working implementation.
- Clear setup instructions.
- Deployment steps.
- Environment configuration.
- A concise explanation of the architecture.
- Links to related documentation.
- A validation process to confirm that the example actually works.

Workstream C — Technical site improvements

Ship targeted improvements from the technical audit, prioritizing issues that directly affect important content.

Potential actions:

- Resolve broken or inaccessible pages.
- Improve rendering of key content.
- Correct indexing or canonical issues where identified.
- Improve internal linking and content discoverability.
- Validate relevant structured data.

- Review crawler access policies based on the intended visibility strategy.

Workstream D — Editorial standards

Develop a lightweight content standard that can be reused across documentation and educational resources.

The standard could define:

- Target audience and technical level.
- Tone and terminology.
- Page structure and heading conventions.
- Answer clarity.
- Code example requirements.
- Accuracy and review process.
- Content ownership and freshness expectations.
- Accessibility requirements.

Deliverable by Day 23: Initial content and technical improvements shipped or submitted for review, with a record of what changed and why.

Phase 4 — Measure expected impact

Timing: Days 20–27

Objective

Evaluate whether the shipped changes correspond with improvements in visibility, content usefulness, and technical performance.

4.1 Before-and-after comparison

For each updated page or query cluster, compare:

- Baseline visibility before the change.
- Visibility after the change.
- Mention rate.
- Citation rate.
- Recommendation rate, where applicable.
- Which URLs are cited.
- Accuracy of the resulting answer.
- Presence and visibility of competitors.

Where possible, use a consistent query set and repeated measurement windows.

Suggested evaluation table

Metric	Purpose
Railway mention rate	Measures brand presence in evaluated answers.
Railway citation rate	Measures how frequently relevant Railway sources are cited.
Citation URL distribution	Identifies which assets are being surfaced.
Recommendation rate	Separates general mention from explicit recommendation.
Answer accuracy	Checks whether Railway's product and capabilities are represented correctly.
Organic impressions & clicks	Measures conventional search discovery for relevant pages.
Qualified engagement	Helps assess whether visitors engage with useful content.
Technical health	Tracks whether identified technical issues were resolved.

Important measurement caveat

Distinguish leading indicators from business outcomes:

- Leading indicators: content coverage, crawlability, page quality, query visibility, citation presence.
- Business outcomes: qualified traffic, developer engagement, relevant conversions, and other agreed growth metrics.

A rise in AI mentions does not automatically mean a rise in qualified users. Likewise, a short-term change in citations does not prove that a single content edit caused the result.

4.2 ROI and learning

Rather than attempting to calculate an exact ROI for every AEO change immediately, document:

- What intervention was shipped.
- Which query or content gap it addressed.
- What changed in the measured results.
- What additional work is justified by the findings.
- What should be deprioritized or tested differently.

This creates a learning loop that can improve future prioritization.

Phase 5 — Iterate and establish the next roadmap

Timing: Days 27–30

Objective

Turn the initial 30-day project into a repeatable process for Railway's growth and content engineering efforts.

Activities

1. Review the results of the initial work.
2. Identify successful and unsuccessful interventions.
3. Update the query evaluation set based on findings.
4. Refine the technical and editorial backlog.
5. Recommend the next 30–60 days of work.
6. Define ongoing measurement and ownership.

Final deliverables

- 30-day results report
- Baseline, work completed, measured changes, limitations, and recommendations.
- Prioritized opportunity backlog
- Remaining content and technical opportunities, with proposed next steps.
- Repeatable measurement process
- Evaluation set, tracking conventions, and a framework for ongoing iteration.
- Next 30–60 day roadmap
- Recommended projects informed by observed results and team priorities.