Generative Engine Optimization

Implementation-ready one-page brief

Bellwether Labs
bellwetherlabs.com

Retrieve it, cite it, reuse it.

One clear identity, clean answers, and evidence a model can quote.

1. Define the entity

- One canonical name, role, and audience across site, docs, and profiles.
- Use @id and sameAs to tie your site, GitHub, and socials together.
- Keep the same wording in About, schema, Wikidata, and descriptions.

2. Make it machine-readable

- Use JSON-LD for Organization, Person, Service/Product, FAQ, Article.
- Set @id, sameAs, author, datePublished, dateModified, and citation.
- Pair markup with the visible page. Don't hide facts behind JS.

3. Publish an AI-readable map

- Add /llms.txt with a short summary, key pages, and .md versions.
- Keep sitemap.xml, robots.txt, and llms.txt consistent and honest.
- Make facts readable without executing JavaScript.

4. Answer real questions

- One question per section, then a 1–3 sentence sourceable answer.
- Use plain headings like 'What is…', 'How…', 'Who it's for'.
- Make the claim and the evidence on the same page.

5. Make evidence citable

- Link to datasets, benchmarks, patents, changelogs, and public work.
- Show author, role, method, date, and what changed.
- Avoid vague claims. One clear fact beats five adjectives.

6. Test like a user

- Pick ten questions people would actually ask about your brand.
- Run them across ChatGPT, Gemini, Claude, Perplexity, and Grok.
- Track citations, reuse, and missing facts. Recheck after changes.

JSON-LD starter

```
{ "@type":"Organization",
  "@id":"https://site.com/#org",
  "name":"Brand",
  "sameAs":[
    "https://github.com/brand",
    "https://www.linkedin.com/company/brand" ] }
```

/llms.txt starter

```
# Brand
> One-sentence summary of what it is and who it's for.

## Docs
- [About](https://site.com/about): what it does and who it's for.
- [Projects](https://site.com/projects): current work.
```