

SEO + AI-readiness + Security Audit

michalismassage.com

Audit · 5 May 2026

Table of contents

SEO + AI-readiness audit · michalismassage.com	1
1. Executive summary	1
2. Technical SEO	2
3. On-page SEO per template	5
4. Structured Data (Schema.org)	6
5. Local SEO	10
6. Content & E-E-A-T	11
7. AI-readiness — the “AI scan”	12
8. Performance & Core Web Vitals	14
9. Internationalisation (NL/EN)	14
10. Conversion & micro-UX	14
11. Security scan	15
12. Prioritised action plan	18
13. Appendix A — measured data per page (33 pages)	19
14. Appendix B — concrete copy-paste fixes	21
15. Closing	22

SEO + AI-readiness audit · michalismassage.com

Date: 5 May 2026 **Scope:** public pages (NL + EN), robots, sitemap, schema, content, AI-discoverability **Method:** automated crawl of 33 pages via curl, head-extraction, JSON-LD parsing, comparison with Google site:-search, performance/security headers **Stack as observed:** Astro (static), self-hosted Lato, Brotli, custom 404, hreflang nl/en/x-default

1. Executive summary

The site is **well built** at its foundation — clear SEO work has been done: titles and descriptions are unique per page and within recommended length, hreflang is correctly implemented, JSON-LD is everywhere, and there is an `llms.txt` (rare — top 3% of the Dutch web). At the same time there are **a handful of silent blockers** that are individually small but together cost noticeable search and AI visibility.

The four highest-leverage fixes (in order):

1. **Fix the schema graph** with `@id` references — there are now multiple unlinked `MassageTherapist/LocalBusiness` entities side by side, with **conflicting** `AggregateRatings` (5.0/600 and 4.9/127 on `/contact`). Google has to guess which entity is your business. One commit solves it.
2. **Add `llms-full.txt`** with actual content (not just links). The existing `llms.txt` is a good index, but AI systems that surface citations (ChatGPT Search, Perplexity, Google AI Overviews) work better with the full text per treatment.

3. **Repair the H1 on Reviews + Quiz** (the Reviews page visually has “Michalis Massage & Sound Healing” as H1 instead of “Reviews”; Quiz has an empty H1 tag in the DOM).
4. **Actively flush old indexed URLs** (/about-me/, /services/, /followup/) — they redirect via 301 but Google keeps them in search results for months without sitemap signal and internal links. Sitemap is fine; just stop linking internally to the old URLs.

Strengths to preserve:

- Full hreflang nl/en + x-default, including self-referential ✓
- Brotli compression + HSTS ✓
- TLS 1.3, with TLS 1.0/1.1 disabled ✓
- No cookies / trackers on first page-load (privacy-friendly) ✓
- Hosted on Vercel: no server-side execution vulnerabilities on the marketing site
- Sitemap-index in place (even though it isn't called sitemap.xml)
- One dedicated page per treatment (10 services), all 1500–2400 words — no thin content
- FAQ page with 18+ questions + correct FAQPage JSON-LD
- A dedicated signature/ page with Service + Offer schema
- llms.txt present and well structured
- 5 active blog posts (all NL) + 9 older posts = reasonable content velocity
- DKIM correctly hooked up via Brevo (b1/b2 records active)

Weaknesses elaborated below:

- No entity graph (@id)
- AggregateRating conflicts
- No AI-bot directives in robots.txt
- No Article schema on blog posts (costly for AI citations)
- No local-content variants for Wageningen/Veenendaal/Bennekom/Renkum
- cache-control: max-age=0 must-revalidate on static HTML — no edge-caching benefit
- 26 of 75 images on the homepage have alt="" — usually correct (decorative), but worth a targeted check
- No Service schema per individual treatment (only on signature)
- The Reviews page shows “655” but the schema reports “600” — Google flags this
- **Security:** only the HSTS header is set; CSP/X-Frame-Options/Referrer-Policy/Permissions-Policy are missing
- **DMARC** is at p=none — monitoring only, no active protection against e-mail spoofing of your domain
- **CAA record** missing — an attacker could force a different CA to issue a rogue certificate

2. Technical SEO

2.1 robots.txt

```
User-agent: *
Allow: /
```

Sitemap: <https://michalismassage.com/sitemap-index.xml>

Verdict: functional, minimal. Returns 200 OK; sitemap pointer is correct.

Recommendation: add explicit AI-bot directives. In 2026 these are the bots that matter for citations in AI answers:

```
# Allow for AI training + citations (= public discoverability via AI tools)
User-agent: GPTBot
Allow: /
```

```
User-agent: OAI-SearchBot
```

Allow: /

User-agent: ChatGPT-User
Allow: /

User-agent: ClaudeBot
Allow: /

User-agent: Claude-Web
Allow: /

User-agent: PerplexityBot
Allow: /

User-agent: Perplexity-User
Allow: /

User-agent: Google-Extended
Allow: /

User-agent: CCBot
Allow: /

User-agent: Applebot-Extended
Allow: /

User-agent: Bytespider
Disallow: /

Why: without these rules, more and more AI systems “play it safe” and skip your site. As a local service provider you want to *be found* via “best message in Ede” inside ChatGPT/Perplexity. Bytespider is TikTok’s crawler — more cost than benefit for your audience.

2.2 Sitemap

- <https://michalismassage.com/sitemap-index.xml> → 200 OK, points to [sitemap-0.xml](https://michalismassage.com/sitemap-0.xml)
- [sitemap-0.xml](https://michalismassage.com/sitemap-0.xml) contains 64 URLs (NL + EN sections cleanly merged)
- <https://michalismassage.com/sitemap.xml> → 404 (custom 404 page)

Recommendation: add a 301 redirect from </sitemap.xml> → </sitemap-index.xml>. Many older tools (incl. legacy Google Search Console submissions) still look at the canonical location. Not critical, but a quick win.

Recommendation 2: add `<lastmod>` to every sitemap entry. Without lastmod Google decides re-crawl frequency itself; with lastmod you prioritise fresh content. Astro can generate this automatically via `@astrojs/sitemap` based on page metadata.

Recommendation 3: split into **logical sub-sitemaps** once you reach ≥ 50 pages: - [sitemap-services.xml](#) - [sitemap-blog.xml](#) - [sitemap-pages.xml](#)

Helps Google Search Console report *which section* has indexation issues.

2.3 Canonicals + hreflang

✓ **Solid implementation.** Per page:

```
<link rel="canonical" href="https://michalismassage.com/message/signature/" />  
<link rel="alternate" hreflang="nl" href="https://michalismassage.com/message/signature/" />
```

```
<link rel="alternate" hreflang="en" href="https://michalismassage.com/en/message/signature-se
<link rel="alternate" hreflang="x-default" href="https://michalismassage.com/message/signatur
```

Things to verify: - Do all EN pages self-canonical to the EN URL? (sample confirms; do a full audit of all 23 EN paths) - Do all hreflang targets actually exist? (one 404 hreflang breaks the whole cluster)

Quick test command (locally in CI):

```
for url in $(curl -s https://michalismassage.com/sitemap-0.xml | grep -oE 'https://[^\<]+'); do
  alt=$(curl -s "$url" | grep -oE 'hreflang="en" href="[^"]+"' | head -1 | sed -E 's/.*/href="')
  if [ -n "$alt" ]; then
    code=$(curl -s -o /dev/null -w "%{http_code}" "$alt")
    [ "$code" != "200" ] && echo "BROKEN $url -> $alt ($code)"
  fi
done
```

2.4 Redirects from old URLs

Google still has these old URLs in its index, currently 301-redirecting:

Old URL	To	Status
/about-me/	/over-mij/	301 ✓
/services/	/message/	301 ✓
/message/thai-massage/	/message/thaise-massage/	301 ✓
/followup/	/	301 ✓

Recommendation: redirects work. To speed up Google’s re-indexation:

1. **Don’t** add the old URLs to the sitemap (correct now)
2. No more internal links to old URLs (sample-check the homepage navigation)
3. Submit the **new** URLs manually via Google Search Console > URL Inspection > “Request indexing”
4. One round of IndexNow pings (Bing/Yandex) — 100 free URLs per day

Other old URLs that might still be lingering: Check Google Search Console → Coverage → “Indexed, not submitted in sitemap”. If anything like ?p=42 (legacy WordPress id), /category/..., /tag/..., /2023/... shows up — return explicit 410s (Gone) instead of 301s.

2.5 Headers + caching

```
content-encoding: br ✓ Brotli
strict-transport-security: max-age=63072000 ✓ HSTS 2 years
cache-control: public, max-age=0, must-revalidate ⚠ no edge cache
```

Recommended cache-control for static HTML:

```
cache-control: public, max-age=0, s-maxage=300, stale-while-revalidate=86400
```

- Browser still revalidates ✓ (no stale HTML in tabs)
- CDN/edge caches for 5 min (s-maxage)
- Edge serves stale up to 24h while it refreshes (stale-while-revalidate)

For **immutable assets** (.webp, .woff2, .css with hash):

```
cache-control: public, max-age=31536000, immutable
```

Headers still missing: - content-security-policy — not SEO, but a security/trust signal -
x-content-type-options: nosniff - referrer-policy: strict-origin-when-cross-origin -
permissions-policy: camera=(), microphone=(), geolocation=()

2.6 404 handling

- Custom 404 page works (/404/ is canonical, NL copy, OG tags fine)
- HTTP status correct 404 ✓
- The 404 page has no noindex — add:

```
<meta name="robots" content="noindex, follow" />
```

(follow so Google follows the navigation links from the 404 to useful destinations)

3. On-page SEO per template

3.1 Title tags

✓ **All 33 titles unique**, lengths 38–86 characters (Google clips around 60 visually but uses everything for ranking).

Issue: two titles above 70 characters may be truncated in SERP: - Thaise Massage: Wat Is Het, Hoe Werkt Het en Wat Kun Je Verwachten? | Michalis Massage (86) - Je Eerste Klankschalen Sessie: Wat Kun Je Verwachten? | Michalis Massage (72)

Recommendation: trim to 50–60 chars and drop the brand suffix when needed: - Thaise Massage: Wat Is Het en Hoe Werkt Het? (44) - Je Eerste Klankschalen Sessie: Wat Verwachten? (47)

3.2 Meta descriptions

✓ All present, 120–170 chars (Google’s sweet spot 140–160).

Slight improvement: descriptions are mostly descriptive without a CTA. Adding “Boek nu”, “Neem contact op”, “Bekijk de agenda” lifts CTR. Example:

- Cupping massage: wat is het, hoe werkt het en wat betekenen die kringen op je huid? Michalis

+ Cupping massage: wat is het, hoe werkt het en wat betekenen die kringen op je huid? Michalis

3.3 H1s

✗ **Two pages with H1 issues:**

/quiz/ — H1 tag present but empty in initial HTML:

```
<h1 class="text-3xl md:text-4xl mb-4" data-astro-cid-tswnd3gl></h1>
```

Likely populated via JS. Google does render JS, but: - Costs crawl budget on every render - Bing/DuckDuckGo/AI bots often render NO JS

Fix: render the H1 server-side. Astro lets you build this statically.

/reviews/ — H1 is Michalis Massage & Sound Healing:

```
<h1 class="text-3xl md:text-4xl mb-4">
  Michalis Massage<br>& Sound Healing
</h1>
```

That’s a brand headline, not a page topic. Keyword relevance for “massage reviews ede” or “klanten-ervaringen massage” is lost.

Fix:

```
<h1>Klantbeoordelingen – 655 reviews voor Michalis Massage</h1>
```

or move the brand block into <p class="eyebrow"> above the H1.

3.4 H2/H3 structure

NL homepage has 9 H2s: 1. Verbind lichaam, geest & ziel 2. Niet zeker welke behandeling bij je past? 3. Diensten & ervaringen 4. Verbeter je sessie 5. Massage in Ede & op locatie 6. Hoe het werkt 7. Wat klanten zeggen 8. Neem contact op 9. Laatste artikelen

Well structured. EN version comparable (9 H2s, equivalent).

Detail: “Verbind lichaam, geest & ziel” is poetic but not keyword-driven. Consider <h2>Massage in Ede: alle stijlen, één plek</h2> as the first H2 below the hero — Google reads early H2s as a topical signal.

3.5 Word counts

Top 5 thinnest pages: | Page | Words | |—|—| | /agenda/ | 377 | | /aanvullingen/ | 550 | | /quiz/ | 597 | | /contact/ | 660 | | /cadeaubon/ | 840 |

Verdict per page: - /agenda/ — rightly compact (event listing), but add a “**Wat zijn ceremonies?**” explainer paragraph (~200 words) at the top for keyword context - /aanvullingen/ — give each add-on (hot stones, cupping, CBD, aromatherapy, red light, theragun) its own 80–150 word section with use-case (“when to choose hot stones”) + duration + extra cost - /quiz/ — interactive, content via JS — fine as long as H1 and intro are static - /cadeaubon/ — add an FAQ block at the bottom (“How long is a gift card valid?”, “Can I pick a custom amount?”, “Can it be personalised?”)

4. Structured Data (Schema.org)

4.1 What’s in place now

Schema type	Found on	Status
MassageTherapist	Every page	✓
PostalAddress	Every page (inside MassageTherapist)	✓
AggregateRating	Every page (often multiple times)	⚠ conflict
LocalBusiness	/contact/, /over-mij/, /massage/signature/	⚠ duplicate entity
OpeningHoursSpecification	/contact/	✓
BreadcrumbList	Sub-pages (not homepage)	✓
FAQPage	/faq/, /massage/signature/	✓
Service	/massage/signature/ only	⚠ missing on other services
Offer	/massage/signature/ only	⚠ same
Person	/over-mij/	✓
Review (individual)	/reviews/	✓
Article / BlogPosting	nowhere on blog posts	✗ big miss
Event	not on /agenda/	✗

4.2 The biggest issue: missing @id graph

What’s wrong now:

Every page has its own MassageTherapist entity with no @id. Google potentially sees **33 different businesses** that happen to share the same name, address and phone. On /contact/ there are also two separate AggregateRating objects (5.0/600 and 4.9/127) — Google has to pick one.

How it should be:

Define **one** Organization/LocalBusiness with a stable @id, and **reference** it from every page. Example global entity (in a shared Astro component):

```
{
  "@context": "https://schema.org",
  "@graph": [
    {
      "@type": ["LocalBusiness", "MassageTherapist"],
      "@id": "https://michalissmassage.com/#business",
      "name": "Michalis Massage & Sound Healing",
      "url": "https://michalissmassage.com",
      "image": "https://michalissmassage.com/images/michalis-portrait.webp",
      "telephone": "+31620159002",
      "email": "info@michalissmassage.com",
      "priceRange": "€60–€200",
      "address": {
        "@type": "PostalAddress",
        "streetAddress": "Kazerne laan 59E",
        "addressLocality": "Ede",
        "postalCode": "6711 RK",
        "addressCountry": "NL"
      },
      "geo": {
        "@type": "GeoCoordinates",
        "latitude": 52.0349,
        "longitude": 5.6840
      },
      "areaServed": [
        { "@type": "City", "name": "Ede" },
        { "@type": "City", "name": "Wageningen" },
        { "@type": "City", "name": "Veenendaal" },
        { "@type": "City", "name": "Arnhem" },
        { "@type": "Country", "name": "Netherlands" }
      ],
      "openingHoursSpecification": [
        {
          "@type": "OpeningHoursSpecification",
          "dayOfWeek": ["Monday", "Tuesday", "Wednesday", "Thursday", "Friday"],
          "opens": "09:00",
          "closes": "20:00"
        },
        {
          "@type": "OpeningHoursSpecification",
          "dayOfWeek": ["Saturday", "Sunday"],
          "opens": "10:00",
          "closes": "18:00"
        }
      ],
      "founder": { "@id": "https://michalissmassage.com/over-mij/#michalis" },
      "aggregateRating": {
        "@type": "AggregateRating",
        "ratingValue": "5.0",
        "reviewCount": "655",
        "bestRating": "5"
      },
      "sameAs": [
        "https://www.facebook.com/michalissmassageandsoundhealing/",

```

```

    "https://www.instagram.com/<account>/",
    "https://www.google.com/maps/place/<place-id>"
  ]
},
{
  "@type": "Person",
  "@id": "https://michalissmessage.com/over-mij/#michalis",
  "name": "Michalis Maroulakis",
  "image": "https://michalissmessage.com/images/michalis-portrait.webp",
  "jobTitle": "Massage Therapist & Sound Healer",
  "worksFor": { "@id": "https://michalissmessage.com/#business" },
  "sameAs": ["https://www.linkedin.com/in/<michalis>/"]
},
{
  "@type": "WebSite",
  "@id": "https://michalissmessage.com/#website",
  "url": "https://michalissmessage.com/",
  "name": "Michalis Massage & Sound Healing",
  "publisher": { "@id": "https://michalissmessage.com/#business" },
  "inLanguage": ["nl-NL", "en"]
}
]
}

```

On sub-pages you don't need to duplicate the full LocalBusiness. Per page only page-specific schemas with @id references:

```

{
  "@context": "https://schema.org",
  "@type": "Service",
  "@id": "https://michalissmessage.com/message/signature/#service",
  "name": "Signature Session",
  "provider": { "@id": "https://michalissmessage.com/#business" },
  "serviceType": "Massage & Sound Healing",
  "areaServed": { "@type": "City", "name": "Ede" },
  "offers": {
    "@type": "Offer",
    "price": "180.00",
    "priceCurrency": "EUR",
    "availability": "https://schema.org/InStock",
    "url": "https://book.michalissmessage.com/signature"
  },
  "aggregateRating": {
    "@type": "AggregateRating",
    "ratingValue": "5.0",
    "reviewCount": "120",
    "itemReviewed": { "@id": "https://michalissmessage.com/message/signature/#service" }
  }
}

```

Important: only add aggregateRating to a Service if you actually have reviews specifically *for that Service* (not just for the business overall). Otherwise keep it on the LocalBusiness level only.

4.3 Resolving AggregateRating conflicts

/contact/ currently has two different ratings side by side (5.0/600 and 4.9/127). Likely one was left over from an older tracking source (Treatwell? Google?) and the other gets updated.

Fix: pick one source of truth. If you have 655 reviews on Google: - Schema rating = Google aggregate (655) - Visible aggregate must **exactly** match (no 600 in code, 655 in copy)

Since 2024 Google has a **strict mismatch policy**: when visible review count \neq schema review count, the entire rich result for the site can be denied.

4.4 Schemas you're still missing

A. BlogPosting on every blog post

Crucial for AI citations. Example:

```
{
  "@type": "BlogPosting",
  "@id": "https://michalissmassage.com/blog/cupping-massage/#post",
  "headline": "Cupping Massage: Wat Is Het en Hoe Werkt Het?",
  "image": "https://michalissmassage.com/images/blog/cupping.webp",
  "datePublished": "2025-09-12",
  "dateModified": "2026-01-04",
  "author": { "@id": "https://michalissmassage.com/over-mij/#michalis" },
  "publisher": { "@id": "https://michalissmassage.com/#business" },
  "mainEntityOfPage": "https://michalissmassage.com/blog/cupping-massage/",
  "wordCount": 2071,
  "inLanguage": "nl-NL",
  "about": [
    { "@type": "Thing", "name": "Cupping massage" },
    { "@type": "Thing", "name": "Bodywork" }
  ]
}
```

B. Event on /agenda/

One Event block per planned ceremony/workshop:

```
{
  "@type": "Event",
  "name": "Cacao Ceremonie & Sound Healing",
  "startDate": "2026-05-24T19:30:00+02:00",
  "endDate": "2026-05-24T22:00:00+02:00",
  "eventStatus": "https://schema.org/EventScheduled",
  "eventAttendanceMode": "https://schema.org/OfflineEventAttendanceMode",
  "location": {
    "@type": "Place",
    "name": "Studio Michalis Massage",
    "address": { "@id": "https://michalissmassage.com/#business" }
  },
  "organizer": { "@id": "https://michalissmassage.com/#business" },
  "offers": {
    "@type": "Offer",
    "url": "https://book.michalissmassage.com/cacao-ceremonie",
    "price": "45.00",
    "priceCurrency": "EUR",
    "availability": "https://schema.org/InStock",
    "validFrom": "2026-04-01"
  },
  "image": "https://michalissmassage.com/images/ceremonie.webp",
  "description": "Yin yoga, sound healing en cacao ceremonie in Ede."
}
```

Events trigger rich results in Google with a date button and can show up in the Google Maps Event tab.

C. Service per treatment

Right now only /signature/ is marked as a Service. Add this to: - Deep tissue (€90–€120) - Sport-massage - Lomi Lomi - Thai Yoga - Shiatsu - Reiki - Duo Massage - Sound Healing

With provider reference back to the central #business id, so the graph stays coherent.

D. HowTo on blog posts where it fits

/blog/klankschalen-sessie-gids/ (“What can you expect”) fits HowTo with steps. AI systems cite HowTo content more often because it provides structural answers.

5. Local SEO

5.1 NAP consistency

Field	Site	Schema	GBP (to verify)
Name	Michalis Massage & Sound Healing	✓	
Address	Kazernelaan 59E, 6711 RK Ede	✓	
Phone	+31 6 20 15 90 02	✓	
Email	(not visible in HTML)	(not in schema)	

Recommendation: add email to the schema. Also add geo.latitude/longitude (Kazernelaan 59E ≈ 52.0349, 5.6840) — improves local-pack rankings.

5.2 Google Business Profile checklist

Not verifiable via crawl, but important: - GBP claimed and verified - Category: primary Massage therapist, secondary Wellness center, Spa, Yoga studio - Opening hours match the site - 20+ photos uploaded (interior, portrait, team, treatment, exterior) - Q&A section enabled, ≥10 question/answer pairs added - Booking button set to book.michalismassage.com - Products/services per treatment (so Google can show them in SERP without a click) - One post per week (event, blog recap, “review of the week”) - Reviews answered actively (within 48h, including 5-star)

5.3 Local content expansion

Competitors within a 15 km radius you’re currently missing: - Wageningen (8 km) — university town, lots of international students/expats - Veenendaal (7 km) - Bennekom (4 km) - Arnhem (15 km) — large

Recommendation: create one page per surrounding town: - /massage/wageningen/ — “Massage in Wageningen — studio 8 minutes by car” - /massage/veenendaal/ - /massage/bennekom/

With: - Own H1 + intro (200 words on travel time, parking, route) - Geo modifier in title+description (“Massage Wageningen”, “Massage in Veenendaal”) - Embedded Google Map with directions - Local notes (“perfect for WUR PhD supervisors looking for an evening session”) - CTA to the booking tool

Note: no pure doorways — add real local value (parking rates, public-transport schedules, local events). Otherwise it’s thin content + manual penalty risk.

6. Content & E-E-A-T

In 2026 **E-E-A-T (Experience, Expertise, Authoritativeness, Trustworthiness)** is a stronger ranking signal than ever for health-adjacent content (massage counts). What you have:

✓ **Experience signals:** - "10+ years of experience" prominently mentioned - Personal photos, name, background on /over-mij/ - 655 reviews

✓ **Trust signals:** - HSTS, no suspicious trackers - Business address + phone visible - Privacy policy + terms-of-service present

△ **Expertise + Authoritativeness under-exposed:** - No author byline on blog posts - No training credentials visible (which massage school? which certification?) - No external authority links to Michalis (interviews, podcasts, guest posts elsewhere)

6.1 Author bylines on blog

Every blog post should have a mini author card:

```
<aside class="author-byline">
  
  <div>
    <p class="byline">Geschreven door <a href="/over-mij/">Michalis Maroulakis</a></p>
    <p class="bio">Massage therapist en sound healer. 10+ jaar ervaring, 600+ 5-sterren review</p>
    <p class="reviewed">Laatst bijgewerkt: <time datetime="2026-01-04">4 januari 2026</time></p>
  </div>
</aside>
```

Google's [Reviewed by]/[Medically reviewed by] pattern is even stronger in 2026. For health/wellness content:

```
<p class="reviewed">Inhoud beoordeeld door Michalis Maroulakis, gecertificeerd massage therapist</p>
```

6.2 Make credentials prominent

On /over-mij/: - Which trainings has Michalis completed (year, institute)? - Member of which professional association? (NGS, VNT, FAGT) - VAT/KvK display (sole-trader signal) - Continuing education / workshops attended?

Also expand the Person schema:

```
{
  "@type": "Person",
  "@id": "https://michalismassage.com/over-mij/#michalis",
  "name": "Michalis Maroulakis",
  "alumniOf": [
    { "@type": "EducationalOrganization", "name": "<massage school name>" }
  ],
  "memberOf": [
    { "@type": "Organization", "name": "<professional association>" }
  ],
  "knowsAbout": [
    "Deep tissue massage",
    "Thai yoga massage",
    "Sound healing",
    "Lomi lomi",
    "Reiki"
  ],
  "hasCredential": {
    "@type": "EducationalOccupationalCredential",
```

```

    "credentialCategory": "Professional certification",
    "name": "<your certification>"
  }
}

```

6.3 Internal linking

The homepage has 41 unique internal links — healthy. Per service page **horizontal linking** to related services is often missing.

Example: /message/deep-tissue/ doesn't currently link explicitly to: - Sportsmassage (similar audience) - Add-ons (deep tissue + cupping combo) - Hot stones add-on - Blog "Massage chronic stress"

Recommendation: add a "Related / pairs well with" section at the bottom of each service page with 3 cross-links + keyword-style anchor text.

Anchor-text examples: - ✓ Sportmassage in Ede - X Lees meer

7. AI-readiness — the "AI scan"

This is where most sites in 2026 leave points on the table. Three concurrent forces are at play:

1. **AI citations** in ChatGPT Search, Perplexity, Google AI Overviews, Claude, Bing Copilot
2. **AI shopping** (driven by Operator, ChatGPT Agents) — booking directly via AI
3. **AI discovery** in voice assistants (Apple Intelligence, Google Gemini)

7.1 llms.txt — already present ✓

GET https://michalismassage.com/llms.txt → 200 OK

Cleanly structured: title + description + sections (Treatments, Group experiences, Practice info, Blog). **Top 3% of the web has this.**

Mini improvements: - Put a **one-sentence elevator pitch** right at the top that LLMs can use in summaries: > **Michalis Massage** is a holistic massage and sound-healing practice in Ede (Veluwe, Netherlands). Active since 2014, 600+ five-star reviews, specialised in Signature Sessions that combine massage with sound healing. - Add a ## Booking section with a direct booking URL and explanation - Add ## Pricing with price ranges per treatment

7.2 llms-full.txt — to add

This is the **recommended** companion to llms.txt: the **full content** of all service pages in markdown, in a single file.

Reason: AI systems that support this (Anthropic Claude, OpenAI GPT, Perplexity) can then understand your whole site in one request rather than 30 separate fetches.

Generate automatically via an Astro build script:

```

// scripts/generate-llms-full.js
import { glob } from 'glob';
import fs from 'fs';
import { remark } from 'remark';
import strip from 'strip-markdown';

const files = await glob('src/pages/**/*.{md,mdx,astro}');
let out = '# Michalis Massage – Full Content\n\n';

```

```

for (const f of files) {
  const content = fs.readFileSync(f, 'utf-8');
  // strip frontmatter, JSX, render to plain markdown
  out += `\n## ${pageTitle(f)}\n\n${plainText(content)}\n\n---\n`;
}
fs.writeFileSync('public/llms-full.txt', out);

```

Real-world examples: Vercel, Anthropic, Cloudflare all publish /llms-full.txt.

7.3 Semantic structure for AI extraction

LLMs parse your site better when you provide:

A. TL;DR blocks at the top of every service page (50 words max):

```

<aside class="tldr" role="note">
  <h2>In het kort</h2>
  <p><strong>Signature Session (180-200 min, €180-€200):</strong>
  combinatie van bodywork, sound healing, aromatherapie en reiki.
  Geschikt voor wie hun lichaam, geest en energie tegelijk wil resetten.
  Boek 5+ dagen vooruit.</p>
</aside>

```

B. Definition blocks for terms that aren't widely known:

```

<dl class="glossary">
  <dt>Lomi Lomi</dt>
  <dd>Hawaiiaanse massage met lange, vloeiende streken die werken op het hele lichaam tegelijk</dd>
  <dt>Sound healing</dt>
  <dd>Gebruik van klankschalen en gongdrums op of rondom het lichaam om vibraties te genereren</dd>
</dl>

```

<dl>/<dt>/<dd> tags are semantically unambiguous; LLMs recognise them as definitions.

C. Q&A in true FAQ form (already done well on /faq/ with FAQPage schema)

7.4 AI-friendly URL structure

Treatment names contain a few “readability” deviations: - /massage/duo-massage/ (NL) - /en/massage/couples-massage/ (EN) - /massage/thaise-massage/ (NL) - /en/massage/thai-massage/ (EN)

✓ Well done.

△ One observation: /sound-healing/ is not under /massage/, and /en/sound-healing/ is also **not** under /en/massage/. Logical, since it isn't massage. However: the NL URL is /sound-healing/ (no localised slug) while /sound-healing-en-klankschalen/ or /klankschalen/ would do more keyword work in the Netherlands. People search “klankschalen Ede” ~5x more often than “sound healing Ede”.

Recommendation: add a 200-OK secondary URL /klankschalen/ that canonicals to /sound-healing/ — or 301 redirect /klankschalen/ → /sound-healing/ and ensure the slug “klankschalen” appears strongly in the H1 + body.

7.5 Content for LLM citation opportunities

LLMs cite content that: 1. Is **list-formatted** (5 reasons, 7 steps) 2. Makes **comparisons** (vs.) 3. Contains **specific** numbers, names, places 4. Has a **source** attached (Michalis Maroulakis, year, certification)

Content ideas that score high in AI:

- “Massage in Ede: 5 stijlen vergeleken (deep tissue vs thai vs lomi lomi)”
 - “Hoeveel kost een professionele massage in 2026? Prijsguids Nederland”
 - “Klankschalen of meditatie: wat werkt beter tegen burnout?”
 - “Massage tijdens zwangerschap: wat mag wel, wat niet”
 - “Cadeaubon massage: wat kost een goede behandeling in Ede en omgeving”
-

8. Performance & Core Web Vitals

Not fully measurable without real-world LCP/INP/CLS data (Chrome UX Report). What I can see:

✓ **Done well:** - Brotli compression - Preload of critical fonts (<link rel="preload" as="font" ...>) - WebP images throughout the site - fetchpriority="high" on the LCP image (at least on /reviews/) - Lato self-hosted (no Google Fonts request — privacy + performance)

△ **Watch points:** - HTML 22 KB uncompressed homepage = ~7 KB gzipped/brotli — fine - 75 images on the homepage — verify all below-the-fold ones have loading="lazy" - No preconnect visible for book.michalissmassage.com — add for a faster booking flow: `html <link rel="preconnect" href="https://book.michalissmassage.com" crossorigin>` - The custom 404 page loads **all** font preloads — unnecessary weight for a 404. Split the 404 page from the main layout.

Recommendation: run [PageSpeed Insights](#) monthly and log scores. Targets: - LCP < 2.0 s - INP < 200 ms - CLS < 0.1

9. Internationalisation (NL/EN)

✓ **Solid:** - /en/ URL prefix for the EN version - Full hreflang nl/en/x-default + self-referential - Menu toggle NL/EN visible - The EN version has its own titles, descriptions, content (not machine-translated)

△ **Two things:**

1. **Inconsistent slug naming** between NL and EN:

- /over-mij/ ↔ /en/about-me/ ✓
- /massage/duo-massage/ ↔ /en/massage/couples-massage/ ✓
- /massage/lomi-lomi/ ↔ /en/massage/lomi-lomi-massage/ △ slugs differ
- /massage/shiatsu/ ↔ /en/massage/shiatsu-massage/ △
- /massage/sportsmassage/ ↔ /en/massage/sportsmassage/ (identical — but note: “sportsmassage” is not natural English; “sports massage” with a space is the idiomatic form)

For consistency, redirect /en/massage/sportsmassage/ → /en/massage/sports-massage/.

2. **Schema inLanguage** is missing. Add to your WebPage schema per page:

```
{ "@type": "WebPage", "inLanguage": "nl-NL", ... }
```

10. Conversion & micro-UX

Not strictly SEO but determines whether organic traffic turns into revenue:

What’s good: - Sticky book button on mobile - Quiz as a micro-conversion - Gift-card CTA prominent around holidays (assumed)

Watch points:

- **Click-to-call** should be above-the-fold on the mobile contact page:

```
<a href="tel:+31620159002" class="cta-call">
  <svg>📞</svg>
  Bel direct: 06 20 15 90 02
</a>
```

- **WhatsApp button** with pre-fill text:

```
<a href="https://wa.me/31620159002?text=Hi%20Michalis%2C%20I%20have%20a%20question%20about%20" class="cta-whatsapp">
  WhatsApp us
</a>
```

- **Reserve button persistent in header** (looks like it is — verify on all viewports)
- **Trust bar just under the hero:** “5.0 ★ · 655 reviews · 10+ years experience” — social proof at the top of the fold

11. Security scan

11.1 TLS / certificate

Item	Finding	Verdict
Certificate authority	Let's Encrypt R12	✓ standard, fine
TLS 1.0 / 1.1	Not supported	✓ good (legacy versions disabled)
TLS 1.2	Supported (ECDHE-RSA-AES128-GCM-SHA256)	✓
TLS 1.3	Supported (TLS_AES_128_GCM_SHA256)	✓
Certificate verify	OK	✓
HSTS	max-age=63072000 (2 years)	✓ good config
HSTS includeSubDomains	not set	⚠ subdomains unprotected
HSTS preload	not set + not preloaded	⚠ can be improved

Recommendation: extend HSTS to:

```
strict-transport-security: max-age=63072000; includeSubDomains; preload
```

and submit the domain at hstspreload.org. **Conditions** for preload: 1. book.michalismassage.com and any other subdomains MUST also be HTTPS-only 2. HSTS is sticky — if you later want to host a subdomain on plain HTTP, you'd have to wait 1 year

11.2 Security headers

What's set:

```
strict-transport-security: max-age=63072000 ✓
server: Vercel
```

What's **missing**: | Header | Status | Risk | |—|—|—| | content-security-policy | ✗ | XSS mitigation absent | | x-content-type-options: nosniff | ✗ | MIME-sniffing attacks | | x-frame-options or frame-ancestors | ✗ | Clickjacking | | referrer-policy | ✗ | Leaks referrer to third parties | | permissions-policy | ✗ | No explicit-deny on camera/mic/geo | | cross-origin-opener-policy | ✗ | Spectre-class isolation | | cross-origin-embedder-policy | ✗ | same |

Recommendation — minimum set for a static marketing site:

In Vercel's `vercel.json` or via `next.config.js`/Astro middleware:

```
{
  "headers": [
    {

```

```

"source": "/(.*)",
"headers": [
  { "key": "Strict-Transport-Security", "value": "max-age=63072000; includeSubDomains;" },
  { "key": "X-Content-Type-Options", "value": "nosniff" },
  { "key": "X-Frame-Options", "value": "SAMEORIGIN" },
  { "key": "Referrer-Policy", "value": "strict-origin-when-cross-origin" },
  { "key": "Permissions-Policy", "value": "camera=(), microphone=(), geolocation=(), pa" },
  { "key": "Cross-Origin-Opener-Policy", "value": "same-origin" },
  { "key": "Content-Security-Policy", "value": "default-src 'self'; img-src 'self' http"
}
]
}
}

```

CSP is the most sensitive one — it has to be tuned to the external services you actually load (Google Maps, GA4, Brevo, booking iframe, etc.). Run CSP in **report-only** mode for one week, watch the console violations, then enforce.

11.3 DNS hardening

Mail (DMARC/SPF/DKIM):

Record	Value	Verdict
MX	mail.michalismassage.com	✓ configured
SPF	v=spf1 a mx ip4:195.211.72.0/22 ip4:141.138.168.0/21 ip6:... ~all	⚠ ~all (softfail)
DKIM brevo1	b1.michalismassage-com.dkim.brevo.com	✓ Brevo connected
DKIM brevo2	b2.michalismassage-com.dkim.brevo.com	✓
DMARC	v=DMARC1; p=none; rua=mailto:rua@dmARC.brevo.com	⚠ p=none = monitoring only

Recommendations:

- Stricter DMARC:** after 2–4 weeks of reading DMARC reports (visible in Brevo), upgrade to:


```
v=DMARC1; p=quarantine; sp=quarantine; pct=100; rua=mailto:rua@dmARC.brevo.com; ruf=mailto:
```

 After another 4 weeks move to p=reject. This stops attackers spoofing your domain — a real risk for a business with Google Reviews because phishing on customers (“review your session here”) is popular.
- SPF ~all → -all:** hard-fail instead of soft-fail, provided you’re certain all legitimate mail comes from mx/a or the listed IP blocks. Brevo sends with DKIM, so no SPF change is needed for Brevo itself.
- CAA record missing:**

```
michalismassage.com. CAA 0 issue "letsencrypt.org"
michalismassage.com. CAA 0 iodef "mailto:admin@michalismassage.com"
```

 Prevents an attacker from forcing a different CA to issue a rogue certificate.
- MTA-STTS + TLS-RPT** for inbox protection:


```
_mta-sts.michalismassage.com. TXT "v=STSV1; id=2026050501"
_smtp._tls.michalismassage.com. TXT "v=TLSPRV1; rua=mailto:tls@michalismassage.com"
```

Plus a `https://mta-sts.michalismassage.com/.well-known/mta-sts.txt` with:

```
version: STSv1
mode: enforce
mx: mail.michalismassage.com
max_age: 86400
```

11.4 Application security

Good:

- ✓ Static site (Astro on Vercel) → no server-side execution vulns on the marketing site itself
- ✓ No cookies on first page-load → privacy-friendly
- ✓ No external analytics trackers in the homepage HTML (no GA4/GTM visible — if it lives elsewhere, verify)
- ✓ No mixed content — all HTTP strings are `xmlns="http://www.w3.org/2000/svg"` (XML namespace, no request)
- ✓ No SRI needed because no external `<script src>` on the homepage
- ✓ Google Maps embed via `https://www.google.com/maps/embed` (official Google iframe, not a 3rd-party tracker)

Watch points:

- ⚠ **Booking subdomain** (`book.michalismassage.com`) is likely an external SaaS (the same audit set would have to be repeated there — different stack, different headers)
- ⚠ **No Subresource Integrity (SRI)** — irrelevant for the Google Maps iframe, but if you later load `<script src>` from a CDN, use:

```
<script src="https://example.com/foo.js"
  integrity="sha384-..."
  crossorigin="anonymous"></script>
```

- ⚠ **Iframe (Google Maps)** — consider `<iframe sandbox="allow-scripts allow-same-origin">` plus `loading="lazy"`. Maps works without most sandbox restrictions.
- ⚠ **No rate limiting visible** on the contact form (Vercel form or fetch API). For a spam target, use:
 - Cloudflare Turnstile (no CAPTCHA, free) or
 - Vercel's [WAF Rate Limiting](#)
 - Fallback: a honeypot field bots fill in but humans don't (a `display:none` field in the form, server-side reject if filled)

11.5 Hosting / supply chain

Item	Finding	Verdict
Hosting	Vercel	✓ secure-by-default, automatic TLS
server header	Reveals "Vercel"	acceptable — Vercel's choice, not exploitable
Brotli compression	✓	
Edge caching	<code>max-age=0 must-revalidate</code>	⚠ improve with <code>s-maxage</code> (see §2.5)
Privacy/GDPR	Privacy policy + terms present	✓
Cookie banner	Not analysed; verify it complies with TCF v2.2 EU requirements if you use GA/FB pixel	

11.6 Reputation / review-fraud risk

Specific to your industry (massage + 600 5-star reviews):

- **Risk:** spear-phishing of customers (“Your session is confirmed / cancel here”). Mitigate via strict DMARC (p=reject) so attackers can’t spoof you.
- **Risk:** review-bombing on Google Maps by competitors. Monitor via [Google Business Profile API](#) — get a notification on every new review and flag suspicious 1-stars matching the “first-time reviewer + no text” pattern.
- **Risk:** look-alike domains (michaliskasage.com, michalimassage.com). One-time:
 - Register common typos (.com + .nl)
 - Or: use a service like [DnsTwist](#) monthly to detect typo domains

11.7 Security score

Based on visible signals (not exhaustive — no pen-test):

Category	Score	Notes
Transport (TLS/HSTS)	8/10	TLS 1.3, HSTS 2 years; missing includeSubDomains + preload
Security headers	3/10	Only HSTS; no CSP, X-Frame-Options, Referrer-Policy etc.
DNS hardening	6/10	SPF + DKIM + DMARC present; missing CAA, MTA-STS, p=quarantine/reject
Application	8/10	Static site, no cookies, no external trackers in HTML
Privacy	8/10	Privacy policy, no cookie set on first hit, Brevo for mail
Average	6.6/10	Solid base, one evening of work to reach 9/10

12. Prioritised action plan

P1 — This week (max. 1 day of work)

1. **Fix the schema graph** with @id + @graph (see §4.2). **Highest impact.**
2. **Resolve the AggregateRating conflict** on /contact/ (pick one source)
3. **Add AI-bot directives** to robots.txt (see §2.1)
4. **Fix Reviews + Quiz H1** server-side (see §3.3)
5. **Match visible review count** with schema (655 everywhere or 600 everywhere)
6. <meta name="robots" content="noindex, follow"> on /404/
7. **301** /sitemap.xml → /sitemap-index.xml
8. **Add security headers** via vercel.json (CSP report-only, X-Content-Type-Options, X-Frame-Options, Referrer-Policy, Permissions-Policy) — see §11.2
9. **Extend HSTS** with includeSubDomains; preload + submit on hstspreload.org — see §11.1
10. **Add a CAA record** for Let’s Encrypt — see §11.3

P2 — This month (1-3 days)

11. **Service + Offer schema** per treatment (8 pages)
12. **BlogPosting schema** on all 14 blog posts
13. **Event schema** on /agenda/ for every planned ceremony/workshop

14. **llms-full.txt** generated via Astro build script (see §7.2)
15. **Author byline** on blog posts (see §6.1)
16. **Credentials on /over-mij/** + extended Person schema (see §6.2)
17. **Cache-control headers** improved (s-maxage + stale-while-revalidate)
18. **Trim title-tag lengths** where >70 (see §3.1)
19. **CTAs in meta descriptions** where missing
20. **Step DMARC up to p=quarantine** (see §11.3) after reading DMARC reports
21. **CSP** from report-only → enforce after one week without violations

P3 — Next quarter (week-level work)

22. **Local pages** for Wageningen/Veenendaal/Bennekom/Arnhem (see §5.3)
23. **Klankschalen secondary URL** or redirect strategy (see §7.4)
24. **Internal-linking audit** + cross-links between services (see §6.3)
25. **HowTo schema** where applicable (see §4.4)
26. **Length expansion** on thin pages /agenda/, /aanvullingen/, /cadeaubon/ (see §3.5)
27. **GBP audit** + 20+ photo update (see §5.2)
28. **Performance baseline** (LCP/INP/CLS) tracked monthly
29. **Hreflang validator** as a CI check
30. **MTA-STX** + **TLS-RPT** records for inbox protection (see §11.3)
31. **Move DMARC to p=reject** after 4 weeks of quarantine without issues
32. **Look-alike domain monitor** monthly (see §11.6)
33. **Rate limiting / Turnstile** on the contact form if spam grows (§11.4)

Ongoing

34. **One blog post every 4-6 weeks** with consistent rhythm
35. **Reply to reviews** within 48h
36. **Sitemap lastmod** updated automatically
37. **Validate schema** via [Schema.org validator](https://validator.schema.org/) on every release
38. **Review DMARC reports** monthly (Brevo dashboard or Postmark/Dmarcian)
39. **Re-check security headers** after every deploy via securityheaders.com

13. Appendix A — measured data per page (33 pages)

URL	Title-len	Desc-len	H1	Words	Schema types
/	49	155	Michalis Massage Ede	2294	MessageTherapist, AggregateRating
/massage/	61	156	Massage in Ede	1751	+ BreadcrumbList
/massage/signature/	62	136	Signature Session in Ede	2110	+ Service, Offer, FAQPage, LocalBusiness (no Service)
/massage/deep-tissue/	45	144	Deep Tissue Massage in Ede	2037	(no Service)
/massage/thaismassage/	51	148	Dynamische Thai Yoga Massage in Ede	1914	(no Service)

URL	Title-len	Desc-len	H1	Words	Schema types
/massage/sportmassage/	38	144	Sportmassage in Ede	2372	(no Service)
/massage/lomi-55 lomi/	55	141	Lomi Lomi Massage in Ede	1844	(no Service)
/massage/shiatsu/	58	140	Shiatsu Massage in Ede	1683	(no Service)
/massage/reiki/	55	144	Reiki Healing in Ede	1749	(no Service)
/massage/duo-52 massage/	52	151	Duo Massage in Ede	2232	(no Service)
/sound-healing/	55	142	Sound Healing in Ede	2324	(no Service)
/ceremonies/	73	153	Cacao Ceremonie & Yin Yoga in Ede	1721	(no Event)
/workshops/	58	148	Massageworkshop in Ede	1648	(no Course)
/over-mij/	43	146	Half Grieks, half Nederlands...	1146	+ Person, LocalBusiness
/contact/	42	138	Neem contact op	660	+ LocalBusiness, OpeningHoursSpec
/locatie/	63	166	Bezoek Michalis Massage in Ede	1501	
/faq/	51	133	Veelgestelde vragen (incorrect)	1627	+ FAQPage (18 Q&A)
/reviews/	48	135		1093	+ Review×20+ (no Event)
/agenda/	58	130	Agenda (empty)	377	
/quiz/	61	121	Massage	597	
/cadeaubon/	43	135	Cadeaubon Ede	840	
/aanvullingen/	53	163	Aanvullingen	550	
/evenementen/	58	155	Onvergetelijke Massages...	2655	
/stoelmassage-op-het-werk/	62	149	Stoelmassage op het Werk	4210	
/blog/	48	125	Kennis & Inspiratie	909	(no ItemList)
/blog/cupping-massage/	64	146	Cupping Massage: Wat Is Het...	2071	(no BlogPosting)
/blog/klankschalen-sessie-gids/	72	171	Je Eerste Klankschalen Sessie...	2296	(no BlogPosting)

URL	Title-len	Desc-len	H1	Words	Schema types
/blog/thaise-message/	86	152	Thaise Message: Wat Is Het...	2126	(no BlogPosting)
/en/	51	161	Michalis Message Ede	2175	
/en/message/	60	163	Message in Ede	1829	
/en/message/signature-session/	68	139	Signature Session in Ede	2213	
/en/about-me/	44	132	Half Greek, Half Dutch...	1152	
/en/contact/	45	145	Get in Touch	680	

14. Appendix B — concrete copy-paste fixes

14.1 Updated robots.txt

```
# Michalis Message robots.txt
# Last updated: 2026-05-05
```

```
User-agent: *
Allow: /
```

```
# AI bots – allowed for citation/discovery
User-agent: GPTBot
Allow: /
```

```
User-agent: OAI-SearchBot
Allow: /
```

```
User-agent: ChatGPT-User
Allow: /
```

```
User-agent: ClaudeBot
Allow: /
```

```
User-agent: Claude-Web
Allow: /
```

```
User-agent: PerplexityBot
Allow: /
```

```
User-agent: Perplexity-User
Allow: /
```

```
User-agent: Google-Extended
Allow: /
```

```
User-agent: Applebot-Extended
Allow: /
```

```
User-agent: CCBot
```

Allow: /

Blocked – more cost than benefit
User-agent: Bytespider
Disallow: /

User-agent: SemrushBot
Disallow: /

User-agent: AhrefsBot
Disallow: /

User-agent: MJ12bot
Disallow: /

Sitemap: <https://michalissmassage.com/sitemap-index.xml>

14.2 Augmented llms.txt opener

Michalis Massage and Sound Healing

> Holistic massage and sound healing practice in Ede, Netherlands (Veluwe region). Run by Mi

At a glance

+
+- ****Location:**** Kazernelaan 59E, 6711 RK Ede (Mauritskazerne, edge of Veluwe national park)
+- ****Reviews:**** 655 five-star reviews on Google
+- ****Languages:**** Dutch and English
+- ****Booking:**** <https://book.michalissmassage.com> (24/7)
+- ****Phone / WhatsApp:**** +31 6 20 15 90 02
+- ****Price range:**** €60 (45-min) to €200 (Signature Session 3hr)
+- ****Specialties:**** Signature Session (massage + sound healing combined), Thai yoga, deep tis
+

Treatments

14.3 Site-wide <head> schema snippet

(To be placed in Astro's BaseLayout.astro, once for the entire domain:)

See §4.2 for the full @graph JSON. Place it before any page-specific <script type="application/ld+json"> blocks.

15. Closing

Infrastructure isn't the problem; **the schema graph and AI discoverability are the multiplier work** on top of what's already there. A well-orchestrated LocalBusiness graph + Service-per-treatment + BlogPosting + Event collected under one entity can show visible effect within 4-8 weeks in:

- AI Overview citations for "massage in Ede", "klankschalen Ede", "cacao ceremonie Veluwe"
- Local-pack rankings for the 3-pack ("Massage in Ede" Maps results)
- Featured snippets for long-tail questions ("how long is a Lomi Lomi massage", "what to wear for a thai massage")
- ChatGPT/Claude/Perplexity citations with a direct booking link

If you'd like me to **execute a subset** as concrete code changes (for example the new schema graph + AI robots), say the word — I'll deliver a PR-ready diff against the Astro codebase as soon as it's available.

— *Audit report based on a public crawl, 5 May 2026*