

18a

PART 2 OF 4 — THE AGENTIC WEB SERIES

The new technical SEO: what actually matters now

The unglamorous stuff that just became load-bearing again.



Hi, I'm Lisa

I've been making websites for over 20 years and run an award-winning web development agency based just outside Bristol in the UK. You can find [our website here](#).

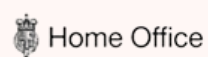
I'm a speaker at events, a writer of all things web, and have taught SEO for longer than I care to remember. I work with clients all over the world, from small one-person-bands to household names. My clients include...

I also founded a not-for-profit to help people find the best web developer for their project, rather than fall foul of a smooth-talking cowboy. From that role, I was asked to represent the industry in discussions with the EU about online privacy reform.

[Instagram](#) [LinkedIn](#)



Sainsbury's



THE AGENTIC WEB SERIES

Four sessions, one theme...

May Why AI keeps choosing the same brands - *replay available*

June **The new technical SEO: what actually matters now - *today***

July The agentic web: is your website usable by AI?

September Selling and building websites in 2027

A QUICK RECAP OF MAY

Where we got to last time

AI recommendation behaves less like loyalty and more like **continuous due diligence** - every time your name comes up, the evidence gets re-checked. It holds up on four signals:

Consistency

The picture of you matches across sources.

Corroboration

Others describe you the way you describe yourself.

Expertise

Demonstrated, not claimed. You show your proof.

Reliability

The site is there, fast, and readable every time.

Missed May? The replay's available - but you don't need it to follow today.

**Technical SEO didn't get less important.
But people stopped spending on it.
And now it's load-bearing again.**

The difference is who's reading your site now. Not just Google deciding where to rank you - retrieval systems deciding whether they can use you at all.

What today is about

- 01 Why the old discipline suddenly carries new weight
- 02 The handful of technical things that genuinely matter now - and why
- 03 The stuff that's new and emerging
- 04 For the agencies: where you prove your value

THE QUESTION UNDERNEATH EVERYTHING

**Can a machine actually read,
understand and trust your site —
right now, today?**

Because if it can't, none of the strategy from May matters.

THE EMBARRASSING ONE

Crawlability

Start here before you do anything else.

You can do everything else perfectly and have it all undone by a single line in your **robots.txt** blocking an AI bot. You don't want to do all this work and discover in six months that your site was blocking ChatGPT's crawler the whole time.

[Free AI readiness audit >](#)

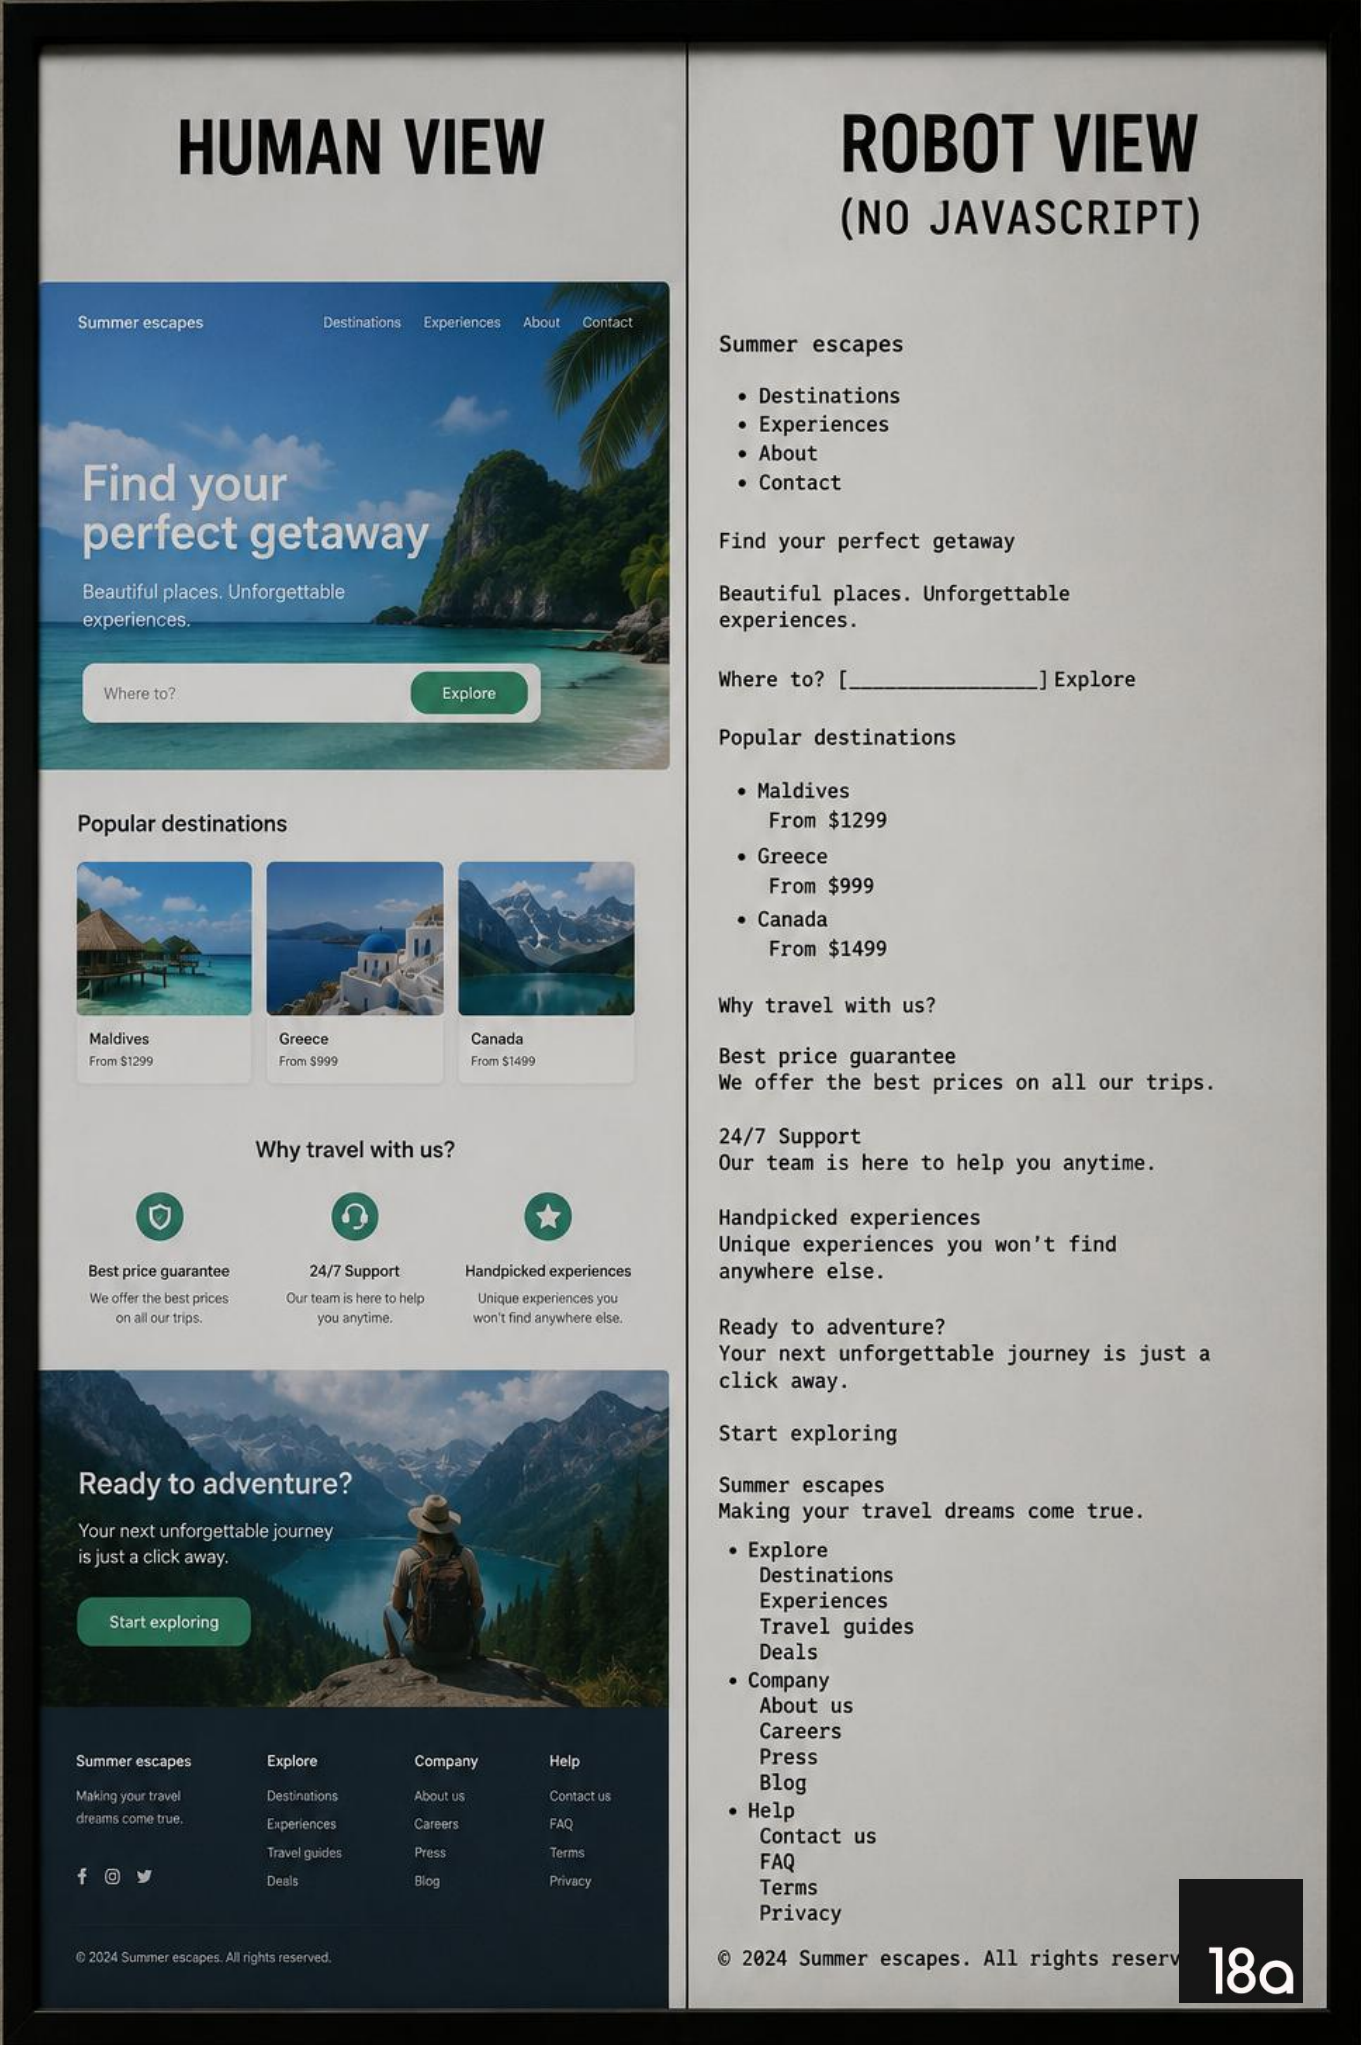


DO BOTS SEE WHAT YOUR VISITORS SEE?

Rendering

A contentious point...

So many modern sites are heavily JavaScript-driven - the content loads *after* the page does. Humans see a lovely page. A bot that doesn't fully render JavaScript might see... almost nothing.



Why rendering matters more for AI than it did for Google

Google, over time

Got fairly good at rendering JavaScript. It will usually wait, fetch and run your scripts before judging the page.

AI retrieval, now

Cost-sensitive and fast-moving. It may not wait around to render a heavy page at all.

If your core content only appears after a pile of JavaScript runs, you're gambling.

THE SAME REASON, SHARPER

Speed


AI is expensive to run. It won't hang around for a sloooow site.

Speed used to be a ranking factor and a UX nicety. Now it's also potentially a **retrieval factor** - a slow site is a less convenient source, so it's a less likely one.



I FEEL THE NEED

The need for speed



HEADINGS, SEMANTICS, CLEAN HTML

Structure and clarity

The boring stuff that's crucial.

A clear heading hierarchy. Semantic HTML - code that says what each bit actually *is*, not just how it looks. This is how a machine works out what your content is about, and which bit is the answer to something. It's the technical foundation under good "chunked" content.

```
<h1>  
  <h2>  
    <h3>  
      <p>  
        <h2>  
          <ul> <li>
```


"When optimising content for AI and humans you must **get to the point early and optimise content to **reduce cognitive load.**"**

— Dan Petrovic

The technical version of "reduce cognitive load" is clean structure. You're making the machine's job easy.

EMERGING — NOT YET PROVEN

llms.txt

A proposed new file that hands AI a clean, efficient summary of your site.

It's an emerging idea - not universally adopted, not proven. But it's cheap to do and worth understanding, so you can make your own call. (AiProfiles generates one, if you want the easy route.)

```
# 18a Productions
> Web development agency, Bristol UK.

## Services
- Web design & build
- Technical SEO

## Key pages
- /services
- /case-studies
```

THE ELEPHANT IN THE ROOM

Structured data

There's been research - Ahrefs and others - suggesting it doesn't strongly affect whether AI cites you.



The FAQ schema myth

WHAT PEOPLE HEARD

"Google stopped showing FAQ rich results - so structured data is dead."

WHAT ACTUALLY HAPPENED

Google dropped the expanded FAQ *boxes in search* .
That's a SERP-display change - not whether schema helps a machine parse your content.

Don't let one story flatten the whole topic.

What that research does and doesn't say

Those studies mostly tested brands that were **already** being indexed and mentioned. They test whether schema gets you cited *more* - not whether it helps AI understand you correctly. Those are two different jobs.


Structured data isn't a discovery tool. It's a clarity tool.

The agencies treating schema as a tick-box - "*we added Organization markup, done*" - are the ones seeing no effect. The value isn't in *having* schema. It's in how richly and accurately you describe things, and how you connect them together. That's what almost nobody's doing well yet.

Google Multi-Stage Pipeline

These studies mostly tested trends that were already being indexed and monitored.

They test whether schema gets you cited more – not whether it helps AI understand you correctly. These are two different jobs.



Not all structured data does the same job

The research that said "schema doesn't move the needle" pooled these two together. They're completely different things.

ENTITY SCHEMA

Identity infrastructure

Tells AI **who you are** - whether you're a resolvable, trustworthy entity or just an ambiguous website that might be anyone.

Organization Person sameAs

This is the half the research never measured.

CONTENT SCHEMA

Page-level labelling

Tells AI **what a page is** - useful tidying, addable any time after build, but not the thing that establishes you.

Article FAQ HowTo

Fine post-build. Not where the value is.

THE NUMBER THAT SHOULD WORRY YOU... OR MAKE YOU FEEL BETTER!

2% of sites are doing what they could with entity schema.
(Our own research, across the sites SiteVitals has scanned.)

[Run free report →](#)

Want to know if you're in the 2% or the 98%? You can [run the same free report on your own site](#).

Entity Schema Reality Check

The state of entity data on the web



Only 2%

of websites have complete entity schema with `sameAs` and `knowsAbout`.



No entity schema at all

26%



Entity schema present, missing both links

51%



Has `sameAs` but missing `knowsAbout`

18%



Complete: entity type + `sameAs` + `knowsAbout`

2%



Why it matters:

Complete entity data helps AI systems understand, trust, and cite your brand.

 sitevitals.co.uk

The 5-Stage AI/SEO Pipeline



1 Crawling

Bots discover URLs.



2 Indexing & Parsing

Content is parsed and entities extracted.



3 Knowledge Graph Construction

Entities are connected to authoritative sources.



4 Retrieval

Relevant, trusted entities are pulled into consideration.



5 Answer Generation

AI generates answers using the retrieved context.



SECTION THREE

New ways of approaching all this...

WHERE IT GETS INTERESTING

Rich, specific markup

Stop labelling *what* a thing is.

Start describing everything you know about it.

A testimonial on your page is just words to a bot - unless you tell it: *this is a review, here's who said it, here's who they are, here's what it's about*. The more you tell it, the more it can do with it.

⚙ The technical bit — a `Review` type, the author as a `Person` / `Organization`, an `itemReviewed` pointing at the specific service, and the author's identity corroborated.

THE "OOH" ONE

The **sameAs** trick

Got a testimonial from an impressive client? Don't just name them. Use **sameAs** on the author to link their identity out to their real LinkedIn profile, their company website - wherever they're verifiably *them* on the web.

You're not just claiming a big name said something nice. You're handing the machine a way to verify the person is real, is who you say, and is genuinely connected to the brand. That's the difference between an unverifiable boast and a corroborated, machine-checkable credential.

Remember the corroboration signal from last month? This is what it looks like when you build it into your site.

```
"author": {  
  "@type": "Person",  
  "name": "Jane Doe",  
  "sameAs": [  
    "linkedin.com/in/janedoe",  
    "janescompany.com",  
    "crunchbase.com/person/jane"  
  ]  
}
```

REAL LIFE CLIENT QUESTION

Categorising testimonials

A client came to me recently saying Gemini had told them to categorise their testimonials. Which sounds odd - until you understand what's underneath it.

PLAIN ENGLISH

Ask AI "is this agency any good for e-commerce?" and a wall of generic five-star quotes is far less useful than one it can tell is about an e-commerce project, from a retail client, about conversion work. Categorising - by service, industry, outcome - lets a machine pull the relevant one instead of treating them as one lump.

⚙️ The technical bit — point `itemReviewed` at the specific `Service / Product`, and use `about / keywords` to tag the theme. So not "5 stars for Agency X" — but "5 stars, for the e-commerce replatforming service, from a retail client, about conversion rate." Now it's retrievable by topic.

The entity brief

Four questions answered *before* a build starts. Not about schema fields - about the thinking schema forces, done early enough to change decisions.

- 1 What are the primary entities?**
The business, the key people, the products, services, locations.
- 2 What claims are we making - and where's the proof?**
On the page, in the markup, across the web. What would a sceptical machine find?
- 3 What external profiles must exist before launch?**
Wikidata, LinkedIn, Crunchbase, industry bodies - corroboration in place on day one.
- 4 What identifiers tie it all together?**
Consistent @id values so every page references the same entity, reliably.

WHY BEFORE, NOT AFTER

Deciding what your testimonials relate to is really deciding what your service pages should be - which decides your navigation.

*Get it right at brief stage and it's free.
Retrofit it later and you're 301-ing URLs
and reshuffling a sitemap that's already
indexed.*

A REAL ONE FROM THIS WEEK

A client's entity data said they'd existed since 2004. Their company registration showed 2023.
(A legal restructure - perfectly innocent. But to a machine doing a background check, the numbers don't add up. That's a red flag against your trustworthiness that nobody would ever think to check.)

The corroboration audit

Checking that the story about your client is consistent everywhere a machine might look - not just on their own site.

1

Does a Wikidata entry exist — and is it right?

It feeds Google's Knowledge Graph. A surprising number of businesses have none, or an out-of-date one.

2

Is there a Knowledge Panel — and is it correct?

If Google's panel says something wrong about you, that's what AI repeats. Know what it says, and how you'd go about fixing it.

3

Is your NAP consistent everywhere?

Name, address, phone - identical across LinkedIn, Companies House and industry listings, not just local directories.

4

Do your sameAs links actually check out?

Every link in your schema should point to a live, accurate profile. Dead or wrong ones undermine the whole picture.

WHAT YOU'RE ACTUALLY BUILDING

Every consistent match the web throws up is a background check passing. Do it enough, in enough places, and your client stops being a "maybe" and becomes a known quantity.

Worth packaging as a named, scoped service - not a tab in an audit spreadsheet nobody acts on.

No authoritative reviews? Do it with your products

The testimonial trick - verify, specify, categorise - works just as well on a product. Same three moves, retail edition.

1 · VERIFY

Prove it's the real thing

Tie the product to its real-world identifiers and link the brand to the manufacturer. The machine can verify you against an authoritative record.

`gtin` `mpn` `brand` `sameAs`

2 · SPECIFY

Describe the offer fully

Price, availability, condition, delivery, returns - modelled, not buried in prose. This is how the systems that represent you understand the deal correctly.

`offers` `availability` `shipping`

3 · CATEGORISE

Make it retrievable

Material, size, colour, specs, use-case - tagged explicitly. So AI can surface YOUR product for a specific query, not lose it in a wall of listings.

`additionalProperty`

WHERE THIS ACTUALLY WORKS

This data does its work upstream — it's how Google's Shopping and indexing systems understand your product before any answer gets written.

But the live check still reads your visible page. So the discipline is making sure your structured data and what's actually on the page agree. When they don't, that's the mismatch that drops you out.

KEEP AN EYE ON IT

Markdown versions of pages

— *and the idea of llms-full.txt*

Stripped-down, "marked-down" plain versions of pages that are more efficient for machines to read. These are future-leaning ideas - not yet recommended. Worth keeping an eye on rather than rushing into.

“You don’t chase rankings. You build content for LLMs to reference, even if your page never breaks into the top 10. The focus shifts from traffic to trust signals: clear formatting, semantic structure, and authoritative insights.”

NEIL PATEL, SEPT ‘25

COMING IN JULY

Where all of this is heading

Clean structure, machine-readable summaries, reliable access - it's all the groundwork for the next shift: sites that AI doesn't just *read*, but **acts inside**.

That's July's talk.



None of this is a one-time job

You can fix all of this on a Tuesday and have it break on Thursday.

A plugin update drops a header. A redirect chain appears. A deploy blocks a bot. An image bloats a page.

Technical SEO isn't a task you complete. It's a state you maintain.



THE QUIET THINGS, CAUGHT LOUD

The ways a website breaks when *nobody's watching*

03:12 · THURSDAY

CSP header dropped after a plugin update

Google starts de-indexing the live site by Monday lunchtime.

22:45 · SUNDAY

Canonical tag now points at staging

You're telling search engines the real page doesn't count.

11:30 · THIS MORNING

Well-optimised title overwritten in a page edit

Nobody knows. Rankings start sliding over four weeks.

06:20 · FRIDAY

A third-party script fails silently

Looks fine to you — but your form submissions now fail.

09:04 · MONDAY

A tiny server config change returns a 500

The home page looks completely broken to every crawler.

02:00 · OVERNIGHT

An image bloats and the page slows to a crawl

Now a less convenient source — so a less likely one.

A site issue caught early is a **five-minute fix . The same issue found by a customer, a client, or your boss is a **very different kind of afternoon**.**

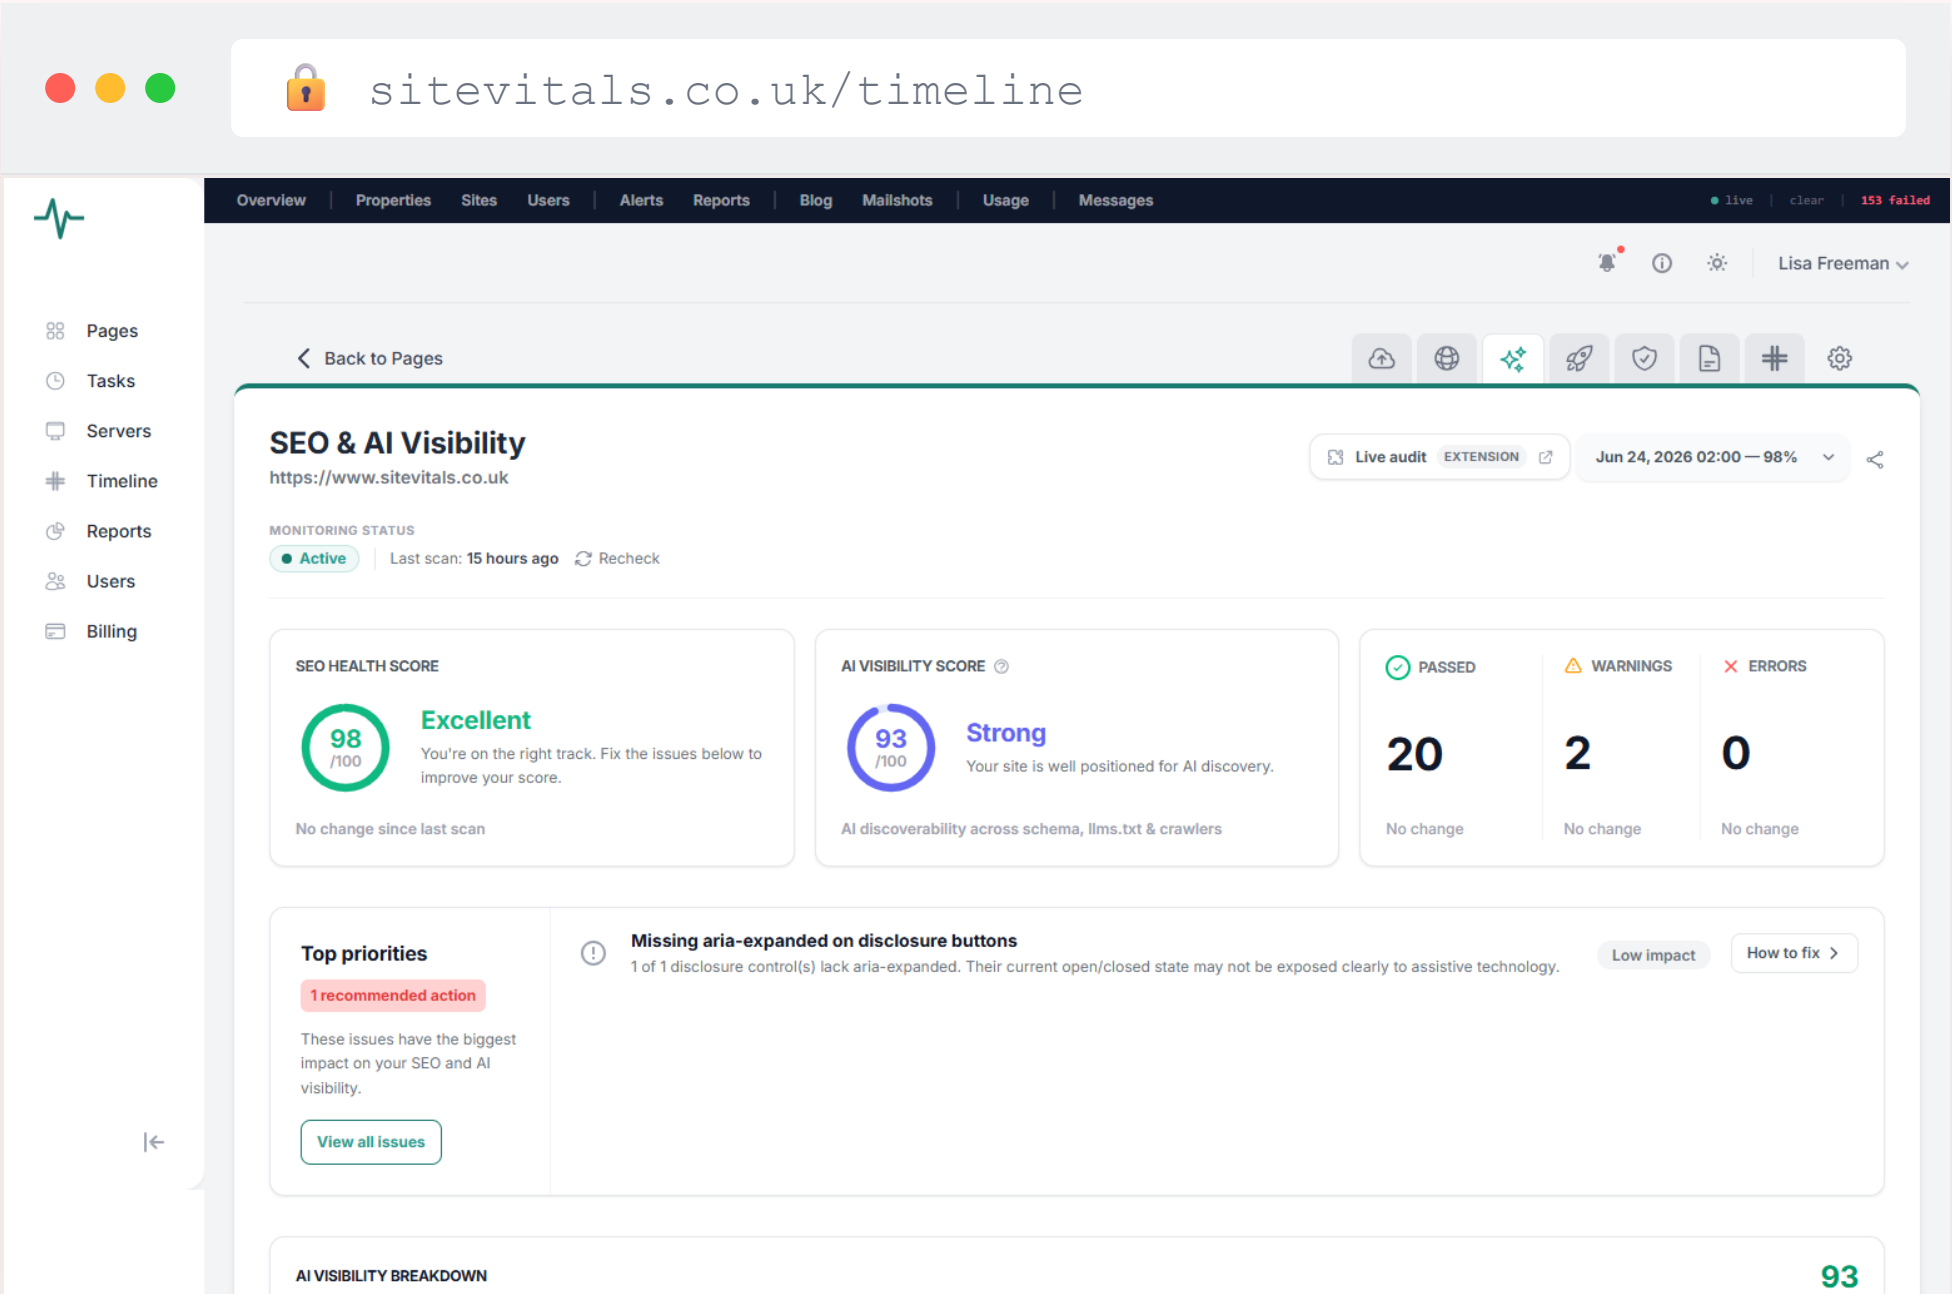
CONTINUOUS SITE HEALTH CHECKS

SiteVitals.co.uk

It watches all the technical stuff we just covered - crawlability, speed, structure changes, broken assets, schema, uptime - and tells you when something dips. So you find out before a customer (or an AI bot) does.

Free account: uptime monitoring, events timeline, and a WordPress plugin.

[Visit SiteVitals →](#)



THE AGENCY ANGLE

What's in it for agencies?

Can AI write your structured data? Yes. Can it write clean semantic HTML? Yes. Can it do it faster than your developer? Yes.



WHERE AGENCIES FIT NOW

**You can't prompt your way to a solution
for a problem you don't know you have.**

The conversation this lets you have with clients

The unglamorous reliability and structure work is now a strategic conversation:

"We know what to look for, and what's actually worth doing about it."

That's a retainer, not a one-off.

That's "trusted advisor", not "vendor".

What changes for in-house marketers

If you're justifying the investment internally, the framing matters. Don't lead with "we need to do technical SEO" - it gets deprioritised every time.

"AI is checking whether our site holds up - and right now, we don't know if it does."

SOMETHING TO TAKE AWAY

**Technical SEO stopped being the luxury bit
the moment machines started deciding
whether you were worth recommending.**

The work you can't see is the work that's holding everything else up.

THE AGENTIC WEB SERIES

Four sessions, one through-line

May Why AI keeps choosing the same brands — *replay available*

June **The new technical SEO: what actually matters now — *today***

July The agentic web: is your website usable by AI?

September Selling and building websites in 2027

PART3 OF 4 - THE AGENTIC WEB SERIES

Is your website usable by AI?

Thursday 16th July, 11am – 12pm



Thank you!

I really hope you found this useful. Any questions - leave a comment on any post on my [Instagram](#) or [LinkedIn](#).

Grab a free [SiteVitals](#) account while you're at it – a complete set of audits on the house.

And I'd love it if you could [leave us a Google review](#) →

