

CHECKLIST

The AI-Ready Website

A working checklist for making a website legible to AI agents and answer engines - the same measures we ship on our own sites.

What AI-readiness actually means

Answer engines, shopping agents and research assistants do not browse your site the way people do. They fetch pages programmatically, extract structure, and quote what they can verify. A site that only looks good to humans is invisible to them; a site that only feeds machines reads as spam. AI-readiness means one well-structured document serving both.

Structure and semantics

- ☐ Exactly one h1 per page; headings form a real outline, not styling
- ☐ Semantic landmarks: header, nav, main, footer - agents use them to find the content region
- ☐ Every page has a canonical URL and matching hreflang pairs for each locale
- ☐ Organization and contact data as schema.org JSON-LD - name, address, contactPoint, sameAs
- ☐ FAQ sections marked up as FAQPage JSON-LD where the content is genuinely Q&A

Machine-readable mirrors

The fastest way to be quotable is to serve the same content in a form that needs no rendering. Three artifacts, all cheap to maintain:

- A Markdown mirror per page (index.md beside index.html) - agents that cannot execute JavaScript still get the full text
- A root llms.txt listing every public page with a one-line description - the emerging convention for AI crawlers
- A sitemap.xml covering every indexable page, regenerated whenever content changes

Everything on this checklist is verifiable: fetch our own pages and you will find llms.txt, .md mirrors, JSON-LD and hreflang pairs exactly as described. Publish nothing you cannot demonstrate.

Content quality gates



Claims point to verifiable evidence - shipped work, live systems, named technologies

uhl-services.ch

- ☐ No duplicate boilerplate between locales; each language is written natively
- ☐ Dates and versions are explicit ('updated October 2026'), not vague recency claims
- ☐ Internal links use descriptive text, not 'click here'

Technical delivery

- ☐ Pages work without JavaScript - agents and crawlers get complete HTML
- ☐ robots.txt is deliberate, not default: allow what should be indexed, nothing more
- ☐ Assets are content-hashed so an agent never quotes a stale version
- ☐ No client-side rendering is required to see primary content

This document is part of the UHL resources collection. Direct download and related work: <https://uhl-services.ch/en/resources/>

