NORDAX DIGITAL

PUBLIC IMPLEMENTATION STANDARD

NDVAS Level 1

Foundational Visibility Architecture

v1.4.1 / TECHNICAL REVISION

The public technical foundation for websites that people, search engines, and AI systems can understand.

Revision date	October 19, 2026
Based on	Level 1 v1.4 / June 2026
Author	Alexz Shepherd, Nordax Digital
Status	Public Distribution Approved - Level 1 only
Public reference	nordaxdigital.com/ndvas
Private standards	Detailed Level 2 and Level 3 requirements remain internal

PURPOSE & INTERPRETATION

The Nordax Digital Visibility Architecture Specification defines technical, structural, and semantic requirements for Nordax websites. It connects search foundations with accessible content and consistent, machine-readable business information.

NDVAS defines implementation and validation requirements. External systems decide what to index, surface, or cite; conformance does not guarantee recognition, rankings, model-training inclusion, or citations.

Level 1: public foundation / Level 2: ongoing maintenance
Level 3: managed identity, integration, and observation

READING THE STANDARD

Scope. Requirements. Evidence.

Level 1 is the public foundation of NDVAS. This revision includes all 20 implementation areas, requirement tiers, context notes, questions and references from the website’s v1.4.1 public specification.

REQUIRED

Nordax acceptance requirements. Record applicability and any reviewed, approved exception.

RECOMMENDED

Adopt when supported by the project’s needs, implementation context, and evidence.

EXPERIMENTAL

Test the consuming system and record limitations before treating an emerging convention as a dependency.

> Technical revision: retains the expanded June 2026 v1.4 scope while correcting performance thresholds, crawler permissions, review / FAQ guidance, identity references, and unsupported outcome claims.

IMPLEMENTATION AREAS

01 Security & trust signals
02 Core metadata architecture
03 Structured data implementation
04 Search & AI crawler policies
05 Crawlability, indexation & crawl health
06 Performance & Core Web Vitals
07 Image & visual search optimization
08 Content structure, semantics & demonstrated expertise
09 Content clusters & topical architecture
10 Answer clarity & generative search readiness
11 Canonical authority control
12 Knowledge graph & entity signals
13 Local & multi-location architecture
14 International SEO
15 Accessibility & inclusive design
16 Digital PR & earned authority
17 AI interpretation & ongoing observation
18 NIST alignment
19 Deployment validation
20 Nordax Digital compliance standard

SECTION 01

Security & trust signals

Objective: Protect visitors and make the website a dependable destination.

REQUIRED

- Enforce HTTPS and redirect alternate hostnames to the chosen canonical hostname.
- Set X-Content-Type-Options: nosniff and an explicit Referrer-Policy.
- Apply an appropriate Content-Security-Policy, using nonces or hashes where needed; test third-party integrations before enforcement.
- Restrict framing through CSP frame-ancestors and, where compatible, X-Frame-Options. Document any legitimate embedding requirements.
- Configure Permissions-Policy for camera, microphone, geolocation, and other capabilities according to actual product needs.
- Enable HSTS after validating HTTPS coverage. Use includeSubDomains and preload only when every relevant subdomain can support the commitment.
- Remove avoidable technology disclosure headers. Monitor for compromised content, deceptive redirects, and broken security controls.
- Keep essential content accessible without intrusive interstitials.

RECOMMENDED

- Review DNS prefetch and external connection hints against the real performance profile.

> These are Nordax implementation requirements. A security header is not an independently guaranteed ranking improvement.

SECTION 02

Core metadata architecture

Objective: Give each indexable page a clear identity and accurate previews.

REQUIRED

- Define the production metadata base URL and a consistent title template.
- Provide a descriptive, unique title, meta description, and canonical URL for each indexable page.
- Identify the business, application, publisher, and language accurately.
- Provide appropriate Open Graph and social-card metadata with real, accessible images and descriptive alternative text.
- Keep publication and modification dates truthful and consistent with visible content.

RECOMMENDED

- Allow large image previews where appropriate with max-image-preview:large. Use 1200-pixel-wide editorial preview images when they suit the placement.
- Validate metadata in rendered HTML and production social previews.

SECTION 03

Structured data implementation

Objective: Describe real businesses and content with accurate, connected records.

REQUIRED

- Publish Organization and WebSite records with stable identifiers, the canonical URL, and supported properties such as name, legalName, inLanguage, publisher, and contactPoint.
- Use LocalBusiness only where applicable, with truthful location, service-area, contact, and opening information.
- Use Article or a suitable subtype on editorial content, with genuine authorship, publication dates, images, and publisher references.
- Provide BreadcrumbList for meaningful page hierarchies. Validate syntax and supported property names at validator.schema.org.
- Ensure markup agrees with content people can see. Test Google-supported rich-result types against the relevant feature guidelines.

RECOMMENDED

- Use Service, Person, ImageObject, and meaningful isPartOf / hasPart relationships where appropriate.
- Use FAQPage for genuine visible questions and answers when useful to a consuming system.
- Use Review or AggregateRating only for an appropriate entity and compliant, genuine, visible review data.

EXPERIMENTAL

- Evaluate Product, VideoObject, Event, DefinedTermSet, HowTo, or Speakable for a real content need and a verified consuming system.

Google retired FAQ rich results in May 2026. FAQPage can still describe a Q&A section; this does not establish an AI citation benefit.
Google excludes self-serving Organization / LocalBusiness reviews from review stars and prohibits aggregating ratings from other websites for that feature.

References [7] [6]

SECTION 04

Search & AI crawler policies

Objective: Make deliberate decisions about search discovery, retrieval, and training.

REQUIRED

- Keep robots.txt reachable and syntactically valid. Include the sitemap location and the intended crawl rules.
- Document separate owner-approved policies for search indexing, AI retrieval, model training, and user-initiated access.
- Allow the search crawlers needed for the agreed discovery strategy, and check CDN / WAF rules against those decisions.
- Verify current provider documentation, agent names, and supported IP verification before creating allow rules. Do not treat a user-agent string as authentication.

RECOMMENDED

- Review Googlebot for Google Search; OAI-SearchBot for ChatGPT search; and provider-documented retrieval crawlers as relevant to the project.
- Record the owner’s independent decisions about training-related controls such as GPTBot and Google-Extended.

EXPERIMENTAL

- Maintain llms.txt, llms-full.txt, or an agent-facing catalog for identified consumers, with ownership and update tests.

OpenAI separates OAI-SearchBot from GPTBot. Allowing training is not a prerequisite for allowing its search crawler.
Google-Extended is separate from Google Search inclusion. User-initiated requests may follow different rules from automatic crawlers.
robots.txt manages cooperative crawling; it is not authentication or a general catalog of entity endpoints.

References [3] [4] [5]

SECTION 05

Crawlability, indexation & crawl health

Objective: Keep public pages discoverable and private systems protected.

REQUIRED

- Generate an up-to-date XML sitemap containing canonical URLs intended for indexing.
- Link important pages through ordinary crawlable navigation and relevant contextual links.
- Use appropriate canonicals, redirects, status codes, and parameter handling. Return a real 404 or 410 when removed content has no relevant replacement.
- Protect staging, administrative, client, and confidential areas with access controls. Use noindex where appropriate for non-private pages excluded from search.
- Monitor meaningful 4xx / 5xx errors and avoid unnecessary redirect chains and loops.

RECOMMENDED

- Use image or video sitemap extensions where relevant. Configure X-Robots-Tag for non-HTML responses when needed.
- Review crawl logs and Search Console in proportion to the size and change rate of the website.

EXPERIMENTAL

- Evaluate IndexNow for a supported search engine and deployment workflow.

> Sitemap submission helps discovery; it does not guarantee crawling or indexing. Do not use robots.txt as a substitute for protecting confidential content.

References [10] [5]

SECTION 06

Performance & Core Web Vitals

Objective: Set measurable performance budgets and preserve usable content.

REQUIRED

- Use LCP under 2.0 seconds as the Nordax project target. Assess INP against 200 ms and CLS against 0.1, using representative devices and field data where available.
- Render essential business and editorial content in server-delivered HTML or an equivalently accessible output; do not make discovery depend on an untested rendering capability.
- Serve properly sized images, reserve their dimensions, and avoid lazy-loading the principal above-the-fold image.
- Control JavaScript and font cost, avoid hydration shifts, and configure caching with a tested invalidation strategy.
- Check mobile layout, keyboard interaction, and essential tasks under realistic network conditions.

RECOMMENDED

- Measure field performance at the 75th percentile where sufficient data exists. Retain lab diagnostics for debugging and prelaunch checks.
- Preload only resources that are truly critical; review third-party scripts and embedded media.

> Google’s published good LCP threshold remains 2.5 seconds. The 2.0-second Nordax target is a stricter internal budget, not a changed Google threshold.

References [8]

SECTION 07

Image & visual search optimization

Objective: Make real imagery useful, accessible, and efficient.

REQUIRED

- Write descriptive alt text for informative images and empty alt text for decorative images.
- Use meaningful filenames where practical, reserve image dimensions or aspect ratios, and preserve the intended composition.
- Serve suitable dimensions and compression, with formats and fallbacks supported by the delivery platform.
- Load below-fold imagery lazily and prioritize the image responsible for the opening experience.

RECOMMENDED

- Connect significant imagery to relevant page copy. Add accurate image metadata and sitemap entries where appropriate.
- Use large, high-quality editorial previews where a distribution feature benefits from them.

> Image format and resolution choices are delivery decisions. No single format stack is a universal AI visibility requirement.

SECTION 08

Content structure, semantics & demonstrated expertise

Objective: Make content clear to people and to machines that process it.

REQUIRED

- Use a clear primary heading, logical heading structure, and appropriate main, article, section, aside, and nav elements.
- Use descriptive link text and relevant internal links.
- Show genuine authorship and accurate dates for editorial material. Attribute claims, quotes, and expertise to the right people.
- Keep service descriptions, business identity, and contact information clear and consistent.

RECOMMENDED

- Use original examples, documented processes, photography, and first-hand evidence where they help the reader.
- Show meaningful update or revision information and answer real customer questions in plain language.

> Content quality is assessed through usefulness and evidence. Formatting, word count, or a fixed update schedule cannot by themselves establish expertise.

SECTION 09

Content clusters & topical architecture

Objective: Connect relevant material so readers can explore a subject.

REQUIRED

- Organize core service or topic areas into useful overview pages with contextual links to supporting material.
- Link supporting pages back to the relevant overview and to related pages when useful.
- Identify and resolve important orphaned pages without manufacturing duplicate or thin content.

RECOMMENDED

- Set content depth and publishing cadence from the audience’s needs, available expertise, and real business priorities.
- Review internal links and navigation depth as the library grows.

> There is no universal required pillar-page word count or number of cluster pages. Topic coverage should follow the subject, not a ranking formula.

SECTION 10

Answer clarity & generative search readiness

Objective: Make supported answers easy to find and understand.

REQUIRED

- Answer the main question directly and supply enough context for the answer to remain accurate when quoted.
- Attribute data and material claims to reliable sources. Distinguish facts, estimates, opinions, and owner-reported results.
- Keep service, location, availability, and identity details consistent with the underlying business.

RECOMMENDED

- Use concise explanations, useful comparisons, steps, tables, and original evidence where they suit the subject.
- Evaluate answer visibility with a documented query set and repeatable observations.

Google specifies no extra technical file or schema requirement for its AI Search features. Eligibility does not guarantee a citation.
Fixed 40-60-word passages, first-200-word rules, and statistics-per-word quotas are not NDVAS compliance requirements.

References [1] [2]

SECTION 11

Canonical authority control

Objective: Maintain a coherent public URL structure.

REQUIRED

- Enforce the intended HTTPS hostname and a consistent policy for duplicate URL variants.
- Provide accurate canonical references for indexable pages; keep internal links and sitemap URLs aligned.
- Map changed URLs to genuinely equivalent destinations. Avoid redirecting unrelated retired content to the homepage.
- Give distinct paginated pages appropriate canonical treatment rather than collapsing every page onto page one.

RECOMMENDED

- Test casing, trailing slashes, tracking parameters, redirect loops, and representative paginated or filtered routes.

SECTION 12

Knowledge graph & entity signals

Objective: Describe the business as a consistent, identifiable entity.

REQUIRED

- Maintain stable identifiers for the organization and relevant services, people, and real locations.
- Keep identity, contact, and service facts consistent across the website and maintained external profiles.
- Use sameAs only for pages that unambiguously identify the same entity; use other suitable links for coverage or references.

RECOMMENDED

- Maintain an accurate entity home and appropriate business profiles.
- Review Knowledge Panel information if a panel exists. Evaluate Wikidata participation only when its notability criteria are met.

> Structured identity supports interpretation; it does not guarantee knowledge-graph inclusion, model training inclusion, a Knowledge Panel, or AI citation.

References [9] [13]

SECTION 13

Local & multi-location architecture

Objective: Represent genuine places and service areas accurately.

Applies to businesses with an eligible physical location or local service area.

REQUIRED

- Provide useful, distinct information for real locations and service areas. Avoid copy-and-paste city pages created only to capture searches.
- Use truthful LocalBusiness details and business-profile eligibility. Do not present a mailing address or virtual office as a staffed, customer-facing location.
- Maintain accurate contact details, services, hours, and location references across relevant owned profiles.

RECOMMENDED

- Build relevant local and industry profiles rather than pursuing a directory-count quota.
- Request genuine feedback in line with each platform’s policies and keep a practical response process.
- Add useful location-specific service details, team information, FAQs, and genuine proof.

> Fresh reviews, profile completeness, and listing consistency should be managed for accuracy and customer usefulness. NDVAS does not assert a universal review cadence or ranking outcome.

SECTION 14

International SEO

Objective: Connect equivalent language and regional experiences.

Applies to projects targeting more than one language or country.

REQUIRED

- Use valid, reciprocal hreflang references, including each page itself, for genuinely equivalent language or regional alternatives.
- Validate alternate URLs, canonical relationships, language / region codes, and the selected HTML, HTTP-header, or sitemap method.
- Keep translated content and business facts accurate, with appropriate regional variation.

RECOMMENDED

- Choose URL structure from business and operational needs. Use x-default where it describes a real fallback.
- Automate alternate-link checks as the number of locales grows; maintain locale-aware identity records.

> Use current validation tools and Search Console capabilities. Do not depend on a retired International Targeting report or an arbitrary sitemap-size threshold.

References [11]

SECTION 15

Accessibility & inclusive design

Objective: Make essential content and interactions usable across access needs.

REQUIRED

- Provide meaningful text alternatives, native controls where suitable, clear labels, and a logical reading and focus order.
- Support keyboard navigation, visible focus, understandable form errors, and accessible dialogs or menus.
- Keep content readable when zoomed or reflowed, respect reduced-motion preferences, and avoid color-only status cues.
- Audit new templates against the agreed accessibility target with automated checks and manual keyboard / assistive-technology review.

RECOMMENDED

- Use WCAG 2.2 Level AA as the target for new work and record any exceptions or remaining issues.
- Use at least 4.5:1 contrast for normal text and 3:1 for large text, with applicable non-text contrast checks.
- Prefer comfortable 44-48 CSS-pixel touch targets where practical. Apply the relevant WCAG minimum size / spacing rules and exceptions.

> A testing score alone does not establish conformance. Record the tested scope, manual checks, and known limitations.

References [12]

SECTION 16

Digital PR & earned authority

Objective: Earn relevant attention through useful work and credible evidence.

REQUIRED

- Avoid manipulative link schemes, fabricated coverage, and undisclosed paid endorsements.
- Attribute research and expert commentary accurately. Distinguish paid placements from editorial decisions.
- Review material link or reputation issues using an evidence-based remediation process.

RECOMMENDED

- Create useful original research, documented case studies, and genuinely helpful expert contributions.
- Track relevant coverage, referral traffic, and business outcomes. Assess editorial quality and audience relevance.

> Domain Rating is a third-party metric, not a Google requirement. No publisher-score cutoff or universal backlink ranking order is specified here.

SECTION 17

AI interpretation & ongoing observation

Objective: Maintain reliable information and evaluate actual discovery behavior.

REQUIRED

- Publish accurate, current descriptions of services and the business in the formats the project supports.
- Keep authorship, source attribution, dates, and important entity references consistent with visible content.
- Document which public endpoints, optional files, integrations, and monitoring surfaces are actually implemented.

RECOMMENDED

- Track brand mentions and linked citations separately, using a recorded query set, platform / mode, locale, and observation date.
- Use observed gaps to improve useful content and identity accuracy; distinguish correlation from causation.

> Optional llms files and integrations can serve identified consumers. They are not a universal prerequisite for search or AI visibility.

References [2]

SECTION 18

NIST alignment

Objective: Make any framework mapping explicit and reviewable.

REQUIRED

- Where NIST alignment is claimed, maintain a written mapping from relevant NDVAS controls to the cited AI RMF guidance, with scope, evidence, and limitations.
- Keep implementation evidence and the responsible owner identifiable.

> NDVAS draws on transparency, traceability, and risk-management principles. NIST has not validated, certified, or endorsed NDVAS. Referencing a framework is not an independent conformance assessment.

References [14]

SECTION 19

Deployment validation

Objective: Confirm the agreed implementation with reproducible evidence.

REQUIRED

- Validate metadata, supported structured data, canonicals, sitemap output, intended crawler policies, HTTPS, and material redirect behavior.
- Check representative page templates and real user journeys on desktop and mobile, including loading, error, and confirmation states.
- Measure the agreed performance budgets and complete accessibility checks for the tested scope.
- Validate locale and real-location requirements where applicable. Check that private material cannot be reached through public routes.
- Record results, release owner, remaining issues, and reviewed exceptions. REQUIRED items must pass or have a documented, approved exception before claiming compliance.

RECOMMENDED

- Retain a dated evidence pack and configure monitoring for meaningful failures after release.

SECTION 20

Nordax Digital compliance standard

Objective: State what conformance covers and who maintains it.

REQUIRED

- Apply the relevant Level 1 controls to Nordax Digital website work. Record non-applicable controls and approved project-specific exceptions.
- Identify the specification version, implementation scope, acceptance evidence, and owner for each compliance claim.
- Review the implementation after material platform, content, or business changes.

NDVAS is a Nordax Digital implementation standard. Conformance describes the tested implementation; rankings, indexing, AI selection, and commercial outcomes remain external decisions.
Level 1 defines the public foundation. Level 2 concerns ongoing maintenance. Level 3 concerns managed entity infrastructure, integrations, and observation. The detailed Level 2 and Level 3 standards remain internal.

PUBLIC REFERENCE

Separate controls. Clear scope.

Representative provider controls from the public specification. Document the owner's decision for each purpose and verify current provider behavior.

Control	Purpose	Policy
Googlebot	Google Search crawling	Agreed public indexing strategy
OAI-SearchBot	ChatGPT search	Agreed search discovery strategy
GPTBot	Potential model-training use	Independent owner decision
Google-Extended	Other Google AI uses	Separate from Search inclusion
ChatGPT-User	User-initiated access	robots.txt rules may not apply

THREE LEVELS / PUBLIC OVERVIEW

Level 1 - Foundational Visibility Architecture

The technical foundation: clear structure, accessible content, consistent identity, and machine-readable business information.

Level 2 - Autonomous Visibility Architecture

Systems that maintain the foundation as your website evolves, keeping visibility signals current when content changes.

Level 3 - Autonomous Authority Infrastructure

Maintained entity records, verified identity information, supported integrations, and observation of how the business is discovered.

> The Level 1 specification is public. The detailed Level 2 and Level 3 implementation standards remain internal. Nordax-managed Level 3 platform integration is delivered through nordax.ai. Implementation scope and evidence determine conformance; recognition, rankings, and citations are decisions made by external systems.

QUESTIONS & ANSWERS

Understanding the standard.

What is NDVAS?

NDVAS is the Nordax Digital Visibility Architecture Specification: our implementation standard for website structure, performance, accessibility, business identity, and machine-readable information.

Who created NDVAS?

NDVAS was created by Alexz Shepherd at Nordax Digital. Level 1 is published as a reference for the websites we build and maintain.

What are the three levels?

Level 1 is Foundational Visibility Architecture. Level 2 is Autonomous Visibility Architecture, covering ongoing maintenance. Level 3 is Autonomous Authority Infrastructure, covering managed entity infrastructure, integrations, and observation.

Which parts are public?

The Level 1 specification is public. The detailed Level 2 and Level 3 implementation standards remain internal. Public summaries explain their scope without publishing their confidential requirements.

Does NDVAS guarantee rankings or AI citations?

No. NDVAS defines implementation and validation requirements. Search engines and AI systems decide what to index, surface, or cite; meeting a technical standard does not guarantee those outcomes.

Is NDVAS endorsed by NIST?

No. NIST has not validated, certified, or endorsed NDVAS. Any NIST alignment claim must be supported by a scoped, documented mapping to the relevant guidance.

PRIMARY REFERENCES

Trace the technical guidance.

Reviewed October 19, 2026. Provider behavior and community policies may change. These references support specific guidance and do not imply endorsement of NDVAS.

[1] Google Search Central: AI features and your website

[2] Google Search Central: Generative AI optimization guidance

[3] OpenAI: Overview of crawlers

[4] Google: Common crawlers and Google-Extended

[5] Google: robots.txt specification

[6] Google: Review snippet structured data guidelines

[7] Google Search Central: Documentation updates, including FAQ retirement

[8] Google: Core Web Vitals

[9] Schema.org: sameAs identity references

[10] Google: Build and submit a sitemap

[11] Google: Localized versions of your pages

[12] W3C: Web Content Accessibility Guidelines 2.2

[13] Wikidata: Notability policy

[14] NIST: AI Risk Management Framework