



SECURITY & COMPLIANCE

Legal and functional description of the service

Intended use, audience, capabilities, dependencies and limitations.

Product	Context
Code	CTX-PUB-06-EN
Document edition	Edition dated 27 July 2026
Effective / review	2026-07-27 / 2027-01-27
Document owner	Niltech Europe S.r.l. · Compliance & Security
Classification	Public

Preamble, nature and effect of this document

Niltech Europe S.r.l., with registered office at Via Calmaggione 5, 31100 Treviso (TV), Italia, VAT No. IT 05614380268 ("Niltech"), adopts this document in order to describe capabilities, workflow, audience, dependencies and limitations without attributing unverified functions or levels.

This document is a corporate transparency and accountability record. It is not a third-party certification, legal opinion addressed to persons other than the company, absolute security warranty or blanket compliance statement; contractual commitments arise solely from the applicable agreements.

This description serves an informational and pre-contractual purpose: it identifies essential characteristics, intended use and material limitations without turning a general representation into a results warranty or a promise of fitness for requirements not expressly agreed.

Personal and material scope

The objective scope includes Context, its public interfaces and processing strictly connected with the described functions. The Context website and service are provided through contextai.legal and use Bluehost infrastructure. OpenAI provides AI processing services; MySQL, Redis, ClamAV, Tesseract, Poppler, ImageMagick and Exim/sendmail are technical components managed within the service environment and are not, for that reason alone, separate subprocessors. Legal-research functions, where enabled by the customer and relevant to the request, may query public institutional sources including Cassazione/SentenzeWeb, Normattiva, AGCM and ANAC; an institutional source does not relieve the user from verifying currency, relevance and the applicable official text.

The relevant operations concern document upload and classification, text extraction and OCR, assisted analysis and synthesis, report drafting, review and research over enabled legal sources. Potential information categories are: authorised accounts, case files, documents, extracted text, OCR results, notes, legal sources, operational logs and assisted outputs. The actual privacy role, lawful basis and extent of processing depend on the contractual relationship and the lawful instructions of the party determining purposes and essential means.

Definitions and interpretation

- "Service" means the Context functions made available under the agreement.
- "Customer" means the legal person or professional entering into the agreement with Niltech.
- "Authorised User" means an individual enabled by the Customer to use the Service under its responsibility.
- "Customer Data" means data, documents, images, instructions and other content submitted or generated on the Customer's behalf.
- "Assisted Output" means a result produced through automated rules or artificial-intelligence components and subject to the stated controls.
- "Further Supplier" means a third party providing Niltech with a technical service relevant to the documented scope.
- "Incident" means an event compromising or capable of compromising confidentiality, integrity, availability, authenticity or resilience.
- "Business Day" means a day other than Saturday, Sunday or an Italian national public holiday.



Specific duties and safeguards

1. Within its assigned role and without prejudice to the Customer's responsibilities, Niltech shall map every function to input, processing, output, user role and control. Performance is governed by necessity, proportionality and traceability, assigned to an accountable person and supported by evidence appropriate to the nature of the risk. Where performance depends on information, authorisation or action by the Customer or a supplier, that dependency is communicated and recorded; absence of the required condition prevents the control from being treated as fully operational.
2. Within its assigned role and without prejudice to the Customer's responsibilities, Niltech shall state quality, format and configuration prerequisites. Performance is governed by necessity, proportionality and traceability, assigned to an accountable person and supported by evidence appropriate to the nature of the risk. Where performance depends on information, authorisation or action by the Customer or a supplier, that dependency is communicated and recorded; absence of the required condition prevents the control from being treated as fully operational.
3. Within its assigned role and without prejudice to the Customer's responsibilities, Niltech shall separate deterministic automation, local tools and AI components. Performance is governed by necessity, proportionality and traceability, assigned to an accountable person and supported by evidence appropriate to the nature of the risk. Where performance depends on information, authorisation or action by the Customer or a supplier, that dependency is communicated and recorded; absence of the required condition prevents the control from being treated as fully operational.
4. Within its assigned role and without prejudice to the Customer's responsibilities, Niltech shall identify where human review is mandatory and what must be checked. Performance is governed by necessity, proportionality and traceability, assigned to an accountable person and supported by evidence appropriate to the nature of the risk. Where performance depends on information, authorisation or action by the Customer or a supplier, that dependency is communicated and recorded; absence of the required condition prevents the control from being treated as fully operational.
5. Within its assigned role and without prejudice to the Customer's responsibilities, Niltech shall document foreseeable errors, unavailability and external dependencies. Performance is governed by necessity, proportionality and traceability, assigned to an accountable person and supported by evidence appropriate to the nature of the risk. Where performance depends on information, authorisation or action by the Customer or a supplier, that dependency is communicated and recorded; absence of the required condition prevents the control from being treated as fully operational.
6. Within its assigned role and without prejudice to the Customer's responsibilities, Niltech shall update description and tests when a material function changes. Performance is governed by necessity, proportionality and traceability, assigned to an accountable person and supported by evidence appropriate to the nature of the risk. Where performance depends on information, authorisation or action by the Customer or a supplier, that dependency is communicated and recorded; absence of the required condition prevents the control from being treated as fully operational.

Purpose and audience

Context is intended for legal professionals and organisations analysing case files, documents and legal sources.

Core functions: document upload and classification, text extraction and OCR, assisted analysis and synthesis, report drafting, review and research over enabled legal sources.

Operational flow

Authorised users enter or collect data, the system organises and processes it, any AI components provide support, and a human reviewer checks it before export or professional use.

Dependencies

Relevant providers and components: Bluehost, OpenAI, MySQL, Redis, ClamAV, Tesseract, Poppler, ImageMagick, Exim/sendmail.

The Context website and service are provided through contextai.legal and use Bluehost infrastructure. OpenAI provides AI processing services; MySQL, Redis, ClamAV, Tesseract, Poppler, ImageMagick and Exim/sendmail are technical components managed within the service environment and are not, for that reason alone, separate subprocessors. Legal-research functions, where enabled by the customer and relevant to the request, may query public institutional sources including Cassazione/SentenzeWeb, Normattiva, AGCM and ANAC; an institutional source does not relieve the user from verifying currency, relevance and the applicable official text.



Limitations

OpenAI supports assisted classification, analysis and synthesis. Tesseract and other local tools support extraction and document revision. Consulted legal sources are identified in the file; the professional verifies authenticity, currency, relevance and completeness before using any result.

Availability, accuracy, supported formats and support depend on contract, input quality and dependencies. The product does not replace a competent professional.

Scope, audience and status of this document

This document is intended for customers, prospects, authorised users, advisers and control functions needing to understand the Context scope. Its specific objective is to describe capabilities, workflow, audience, dependencies and limitations without attributing unverified functions or levels. It applies to the stated document revision and date and must be read with the applicable agreement, order, DPA, technical specifications and controlled procedures.

The Context website and service are provided through contextai.legal and use Bluehost infrastructure. OpenAI provides AI processing services; MySQL, Redis, ClamAV, Tesseract, Poppler, ImageMagick and Exim/sendmail are technical components managed within the service environment and are not, for that reason alone, separate subprocessors. Legal-research functions, where enabled by the customer and relevant to the request, may query public institutional sources including Cassazione/SentenzeWeb, Normattiva, AGCM and ANAC; an institutional source does not relieve the user from verifying currency, relevance and the applicable official text.

Executed agreements and actually approved configurations prevail in case of inconsistency. Public information describes the control programme; it does not turn optional provider capabilities into Niltech controls or automatically attest legal applicability or satisfaction.

Operational definitions

The following definitions support consistent interpretation. Terms defined by law or contract retain the meaning assigned by the relevant source.

- Input: data, documents, images, instructions or parameters supplied to the system.
- Assisted output: result generated or enriched by automated components and subject to review.
- Use case: combination of users, purpose, data, functions and supported decisions.
- Known limitation: condition that may reduce accuracy, completeness, availability or suitability.
- Human review: substantive check by a competent person before professional use.

Requirements and operating process

The process must remain proportionate to the use case, data, exposed persons and possible impact. Relevant Context functions include document upload and classification, text extraction and OCR, assisted analysis and synthesis, report drafting, review and research over enabled legal sources. Every phase must have an owner, verifiable inputs, completion criteria, evidence and an escalation path.

Before activation or publication, the owner verifies that requirements are translated into concrete controls and that controls actually apply to the environment. A merely planned control, an unconfigured capability or a general supplier statement is recorded as a gap, not as an operating measure.

- map every function to input, processing, output, user role and control
- state quality, format and configuration prerequisites
- separate deterministic automation, local tools and AI components
- identify where human review is mandatory and what must be checked
- document foreseeable errors, unavailability and external dependencies
- update description and tests when a material function changes

Roles and responsibilities

Niltech maintains this document, configurations under its control and links to evidence. The customer determines purposes, users, data, authorisations and professional decisions within its scope. Providers are responsible for components defined by their contracts. The primary product audience is legal professionals and organisations analysing case files, documents and legal sources.



No role is assigned implicitly. Product, security, privacy, continuity and AI owners must be identified in the controlled register. Where one person performs multiple roles, material risk, exception or result approval must include proportionate independent review.

- Management: approves residual risk, material exceptions and priorities.
- Owner: maintains requirements, controls, evidence, gaps and reviews.
- Operators: follow instructions, verify inputs/outputs and report anomalies.
- Customer: ensures lawfulness, authorisations, data quality and professional use.
- Suppliers: provide services and information under their agreements.

Evidence, verification and accountability

Every material statement must trace to current evidence: configuration, contract, report, test, ticket, minutes, log or register. Evidence is classified, restricted to those with a need and retained under the applicable schedule. Hash and document revision identify documents but do not alone prove control effectiveness.

Verification assesses evidence existence, relevance, coverage, date, result and limitations. Samples and tests must represent the stated risk. Negative or partial results create an action with owner and due date; they are not omitted from the summary where they affect risk understanding.

- functional specification and approved use cases
- component, dependency and document revision inventory
- quality, accuracy, security and usability tests
- known limitations, issues and distribution records
- human-review and export evidence

Cross-cutting safeguards

Privacy, security, quality and continuity are considered at every phase, not only at completion. Data is limited to what is necessary; access and authorisation follow least privilege; changes and administrative activity are traced proportionately; inputs, transformations and outputs retain provenance sufficient for professional review.

Special-category data, highly confidential information, minors or vulnerable persons require a specific assessment before processing. Where a required control is unavailable or unproven, scope is reduced, a manual mode is used, the data is removed or the use case is suspended.

- minimisation of data, free-text fields, prompts, logs and exports
- named access, role segregation and timely revocation
- format, integrity, completeness and provenance validation
- evidence-based protection of transport, working copies and sharing channels
- degraded mode, escalation and communication where dependencies or controls are unavailable

Exceptions, non-conformity and escalation

A deviation is not accepted by custom. The owner records the affected requirement, cause, impact, exposed data and persons, compensating measures, approver, expiry and closure criterion. The exception is reviewed if risk changes or a measure does not work as expected.

Incidents, possible unlawful processing, loss of data control, outputs with severe impact, contractual breaches, unapproved suppliers or unreliable evidence must be escalated without delay. Current Legal and functional scope: The documented scope includes the website, application service, technical components and external sources used under the agreement.

- contain risk and suspend the affected phase where needed
- preserve evidence, timing, decisions and communications
- involve privacy, security, product, legal or management owners as appropriate
- resume only after measure verification and documented authorisation



Review, change and improvement

The document is reviewed at least every six months and earlier when purpose, audience, data, GDPR or AI Act role, supplier, model, architecture, location, contractual terms or legal requirements change. Incidents, complaints, failed tests and new vulnerabilities trigger an extraordinary review.

Each review records inputs, participants, decision, changes, superseded evidence, remaining gaps and next date. Material corrections are published without retroactively altering the prior document revision. Contact and requests: info@nil-tech.net.

- check change register and related documents
- retest affected controls
- update manifest, PDF, HTML and hashes
- notify recipients where the change affects their rights or duties

Allocation of responsibility and reliance limitations

Within its sphere of responsibility, the Customer warrants the lawfulness of submitted data and instructions, user authorisation, suitable lawful bases and notices, and professional verification of outputs. Niltech remains responsible for activities directly under its control and does not assume the Customer's regulatory, professional or decision-making functions.

Outputs from Context are auxiliary. Unless expressly agreed and subject to mandatory law, they are not legal advice, an expert determination, insurance decision, liability finding, credit assessment or other reserved professional act. The recipient must examine sources, completeness, consistency and consequences before use.

Nothing excludes liability that cannot lawfully be excluded. Outside those cases, attribution, remedies, limitations and quantification principles follow the applicable agreement, taking account of contributory conduct, mitigation duties and foreseeability under the governing law.

Evidence, review, requests and governing law

Every material assertion must be traceable to a contract, approved configuration, register, minutes, test, log or other reliable evidence. Supplier statements and Niltech controls are kept distinct. Absence of incidents is not, by itself, proof that a measure is effective.

Revisions are dated, reasoned and approved. A later revision does not retroactively alter facts or commitments applicable to earlier periods. Published copies are identified by code, date and cryptographic digest; those elements evidence copy integrity, not the substantive effectiveness of described controls.

Reports, clarification requests, rights requests and complaints may be sent to info@nil-tech.net. Niltech verifies identity and authority where necessary, records the request, responds within applicable periods and communicates any reasoned extension or refusal.

Unless mandatory law or a written agreement provides otherwise, Italian law governs interpretation. The Italian text is controlling; the English translation is provided for convenience.

Official sources and legal date

References checked on 2026-07-27. Applicability, role and any relevant commencement dates must be reassessed against the concrete facts.

- Italian Legislative Decree No. 70 of 9 April 2003 — Electronic commerce
- Italian Legislative Decree No. 206 of 6 September 2005 — Consumer Code
- Regulation (EU) 2024/1689 (AI Act) and Commission framework

Document control: the copy identified in the manifest through its SHA-256 digest is the document edition published on the stated date. The header seal is neutral corporate artwork and is not a digital signature, third-party attestation or certification.