LongTermCapabilities Security Overview

Public trust and operating boundaries for the website and principal-led delivery model. This is not a certification, attestation, penetration test, legal opinion, insurance representation, or opportunity-specific control assessment.

Security overview

The public site is static-first, uses local assets and search, does not configure external analytics or a third-party form processor, and does not request sensitive technical material. Client security requirements are qualified and scoped before access.

- Least authority and role separation
- Secure credential handling through approved client channels
- Version control, review gates, logging, and rollback where scoped
- Environment and data-boundary separation
- Data minimization and dependency review
- No secrets committed to public source or client-side code

Data-handling boundaries

Public channels are limited to public-safe qualification. Confidential or regulated evidence moves only after responsibility, scope, contractual terms, and an approved transfer method are clear.

- Public-safe: organization, role, public notice links, high-level system context, timing, and budget range
- Confidential business information: accepted only through an approved private channel after qualification
- Regulated or sensitive information: requires explicit contractual and technical handling decisions
- Prohibited through public forms and email: credentials, private source code, customer records, PHI, financial account data, classified, CUI, export-controlled, source-selection-sensitive, or privileged material

AI use and human authority

AI output is treated as proposed work. Consequential action requires explicit human authority, and generated claims require source support or an unknown state.

- Client data is not placed in unapproved tools
- Data, retention, and model-use boundaries are defined for the engagement
- Human review remains explicit for consequential work
- AI use is disclosed when materially relevant
- Unknowns are not converted into certainty
- Approval and execution are separated where risk warrants

Secure software practices

Engagements may use secure-development concepts appropriate to the system, including least authority, reviewable change, dependency management, versioned evidence, test gates, rollback, and explicit operational ownership.

- Source and change control
- Dependency and configuration review
- Secrets excluded from source
- Review and release gates
- Logging, observability, and rollback design
- Client-owned artifact export and handoff

Business continuity and exit

The delivery model favors client-owned artifacts, exportable formats, documented dependencies, knowledge transfer, and a planned handoff rather than manufactured lock-in.

- Dependencies and external services are documented
- Decision and risk records are exportable
- Handoff and exit expectations are named in scope
- Additional personnel are disclosed and scoped

Framework alignment

NIST AI Risk Management Framework - Engineering alignment

Use the voluntary framework as a risk-management vocabulary for ownership, context, measurement, monitoring, response, and documentation.

Boundary: LongTermCapabilities does not claim NIST certification or universal conformance.

NIST Secure Software Development Framework - Practice reference

Use secure-development concepts to structure software lifecycle, review, dependency, release, and acquisition conversations.

Boundary: Framework reference does not constitute certification, attestation, or a security guarantee.

Section 508 and accessible ICT guidance - Accessibility-aware delivery

Use official accessibility guidance to inform design, development, testing, and procurement evidence when public-sector requirements apply.

Boundary: No Section 508 or WCAG certification is asserted.

ISO/IEC 42001 readiness concepts - Readiness support only

Organize technical inventories, system descriptions, human oversight, evaluation evidence, and change records that may support a qualified management-system program.

Boundary: LongTermCapabilities is not a certification body and does not issue ISO/IEC 42001 certificates.

Least-authority and Zero Trust principles - Architecture principle

Use explicit identity, authorization, segmentation, data minimization, verification, and separation of approval from execution where appropriate.

Boundary: No formal Zero Trust compliance or government approval is asserted.

Responsible disclosure

Report public-site security concerns to consulting@longtermcapabilities.com with the subject 'Security report'. Include the affected public URL, observed behavior, date, browser or client, and a non-destructive reproduction description.

- Do not access, alter, destroy, or exfiltrate data
- Do not degrade availability or conduct denial-of-service testing
- Do not use social engineering, credential attacks, or destructive automation
- No bounty or safe-harbor promise is offered unless separately agreed in writing

Procurement and private document requests

Legal entity details, W-9, insurance certificates, registration records, security materials, standard agreements, and other opportunity-specific documents are shared only through an approved qualification channel when verified and relevant.

Last reviewed 2026-07-23. Public Trust Center: https://longtermcapabilities.com/trust/